Probabilistic Modeling of Grasping Affordances in Robots using Bayesian Networks

Alice Ellmer

Supervisors:
Dr. Giorgio Metta
Dept. of Robotics, Brain and Cognitive Sciences
Italian Institute of Technology

Dr. Joachim Hertzberg
Institute of Computer Science
University of Osnabrueck

A Thesis submitted in partial fulfillment for the degree of
Bachelor of Science in Cognitive Science

Institute of Cognitive Science
University of Osnabrueck
September 2009

Contents

List of Figures

List of Tables

Abstract

Enabling robots to act in unstructured environments is a dauntingly complex task. Approaches based on affordances, the perceived action possibilities offered by the environment, promise to be an important stepping stone. A recent example of the successful application of the affordance based approach in the domain of dexterous manipulation is given by Montesano et al. [1]. Their work builds on the idea of using Bayesian Networks to formalize affordances. These Bayesian Networks are learned to encode relationships between objects, actions and observed effects to one another. In related work, Metta and colleagues described a desirable affordance representation [2] to be probability distributions over action primitives given an object. The current work first aims at replicating the positive results obtained by Montesano et al. in the domain of grasping actions. Then, the problem how the obtained Bayesian Networks can be used to derive probability distributions over action primitives is addressed. It could be shown that in general Bayesian Networks are a suitable formalism for deriving affordances from experience. However, how well these affordances reflect the acquired experience proved to depend heavily on the topology of the network. The results suggest that the taken method generalizes to more complex tasks, but also that it needs to be complemented with mechanisms for inducing goal-oriented behavior in robots.

Keywords: Affordances, Bayesian Networks, Dexterous Manipulation, Robotics

Acknowledgements

I am very grateful to a lot of people, be it for their academic assistance, motivation efforts, discussions, criticisms or simply patience in enduring my questions throughout my work on this thesis. First and foremost, I would like to thank my adviser Dr. Giorgio Metta for giving me the possibility to work at the Italian Institute of Technology in Genoa with the robotic platform iCub. Without his support and that of many more people at the IIT, notably Francesco Nori and Ugo Pattacini, this thesis would have been impossible to realize. Also, without the feedback from Dr. Joachim Hertzberg, this thesis would have lacked much of its clarity, and I am thankful for his experience and clear-sightedness he generously shared with me. Particularly helpful when I was caught up fiddling with implementation issues or nerve-wrecking errors in reasoning have been Dr. Luis Montesano, Dr. Ebert Uphoff and Dr. Stefan Evert.

I would also like to mention that the idea for this thesis originated from my work with Klaus Libertus and Dr. Amy Needham at the Infant Perception Lab at Duke University. In retrospect, I can call it no less than awe-inspiring to see how infants succeed to interact with their environment. When put in this context, the current thesis can not be called much more than an initial groping in the dark in the domain of developmental robotics.

Many thanks also go to Dr. Britta Wrede, Stefan Krüger and Dr. Sebastian Wrede from the Research Institute for Cognition and Robotics for welcoming me at Bielefeld University.

On a more personal note, I thank Nihshanka and Andrej. Nihshanka for encouraging me to try things I was first quite afraid of, and Andrej for helping me through a couple of downs on the way.

Structure of this Thesis

The first chapter of this thesis provides the reader with background knowledge about affordances as they are investigated in ecological psychology and neuroscience. Following this is a chapter about the basic concepts and methods of Bayesian Networks, which may be skipped by the reader familiar with these topics. Chapter three retraces the origins of this work and is essential for understanding and interpreting the rest of this thesis. Chapter four details data acquisition and the process of arriving at Bayesian Networks from experience. Finally, chapter five addresses the issue of how well the Bayesian Networks obtained in chapter four perform in the task of inferring the stability of an action. Chapter six discusses how to arrive at the aforementioned affordance representations as probability distributions. Eventually, chapter seven concludes with a general discussion of the results and possible future directions for research.

1 Introduction & Motivation

(a) Handle for pulling the cabinet

(b) Hollow for opening the drawer

Figure 1: File Drawer

The file cabinet shown in Figure 1 looks like the handle in the left image is supposed to be pulled in order to open the cabinet's topmost drawer. What happens instead when you pull the handle is that you move the whole cabinet itself, which is placed on rollers, rather than only opening the drawer. Figure 1b shows how to actually open the drawer by reaching into a hollow. Bluntly speaking, this can be considered counterintuitive design: the handle highlighted in Figure 1a affords being pulled in order to achieve a certain effect, namely opening the drawer.

Another example of a similarly poor design is shown in Figure 2. The design of this water station is reminiscent of an ordinary old-fashioned water pump: the handle of such a water pump needs to be pulled up and pushed back down in order to pump water through the faucet. However, this is not how this particular water station works: as a modernized version, the lever simply serves as an ON/OFF switch: the lever being in a vertical position (see Figure 2a) meaning OFF and no water flow, and the lever being in a horizontal position meaning ON (Figure 2b) and continuous water flow. Thus, we can say that the lever affords a pumping like action, although it was intended for a completely different action, namely being used as an ON/OFF switch.

(a) Lever in OFF position

(b) Lever in ON position

Figure 2: Water pump

These two examples[1] serve to illustrate what is usually meant when we talk about affordances: the functional significance of entities such as surfaces and objects as seen relative to an actor and her action capabilities. Commonly, affordances are defined to encompass the effects of an action/object pairing, such as the water flow in the example of the water pump. Psychological research strongly suggests that it is indeed the case that we humans are able to almost instantaneously associate a set of suitable actions to objects in our daily lives. Although easily taken for granted, this is a remarkable capability.

[1] Many thanks for the permission to use these images go to Mike Darnell from http://www.baddesigns.com.

The first scientist to explicitly name this phenomenon and to introduce it to a broader audience was James J. Gibson [3, 4]. Although Gibson himself was a psychologist, his ideas were incorporated into the areas of philosophy of mind, neurobiology and artificial intelligence research. The following sections serve to give an impression of the status affordances have in the areas of ecological psychology and neuroscience.

1.1 Affordances in Ecological Psychology

James J. Gibson conceived himself a psychologist, and his ideas naturally had the largest impact in his own discipline. In fact, he and his wife Eleanor Gibson lay much of the foundation for the relatively young discipline of ecological psychology. Research in ecological psychology emphasizes the interaction between the agent's perception and action capabilities. A detailed review of research addressing the developmental aspects of ecological psychology is beyond the scope of this work, and has already been done in excellent form in [5].

How do the perceived properties of the environment interact with the physique and capabilities of an organism living in it? Far from providing a definite answer, yet a famous effort towards it, is a series of experiments on stair climbing affordances conducted by William H. Warren in 1984 [6]. Warren studied the relationships between leg length and optimal and critical riser height[2] of steps in a psychophysical setup. He found the ratio between leg length and riser height to reliably predict optimal as well as critical points in the stair climbing task. Furthermore, these points could be readily visually identified by the participants. This was first evidence for the hypothesis that agents indeed perceive their environment in relation to their physique and motor capabilities. This result is even more astonishing since the relationship between the agent and its environment was quantified only through a simple ratio between environmental and bodily measures. Similar results were obtained regarding perceptual judgements and motor decisions in the domain of manual interactions [7].

Furthermore, evidence was collected to support the general view point that actions and the processing of sensory input are closely related. For example, Flanagan [8] demonstrated that the gaze patterns during executing and observing a block stacking task are highly similar. He concluded that the action plans used throughout such tasks are also used while observing and understanding other people's actions.

In 1997, van der Meer [9] conducted a fascinating study showing that interaction between sensory feedback and motor decisions is already present shortly after birth. In this study, newborns were placed on their back in a dark room. Above their chest

[2] The most comfortable and economical height of a step for a given subject was termed "optimal riser height". The maximum step height a subject could ascend was termed "critical riser height".

a light beam was projected, such that the baby was in complete darkness unless it raised its hand and thus brought her hand into the light beam. Analysis showed that the motor activity of the newborns changed systematically depending on the position of the light beam. These results suggest that the babies aimed at increasing the portion of time they saw their arm, i.e. received visual feedback from their motor actions. Van der Meer also observed that the babies reduced the velocity of their arm before entering the light beam, thus possibly anticipating the visual feedback which would be provided shortly after. This remarkable behavior shows that even the seemingly unstructured neonatal movements can be considered as serving the purpose of visuomotor coordination.

For the interested reader, Von Hofsten provides [10] a more complete review of the interplay between motor and cognitive development in infancy.

1.2 Affordances in Neuroscience

Much of the work on affordances in neuroscience based on electrophysiological studies performed with macaque monkeys. Figure 3 shows a side view of the cerebrum of an adult macaque monkey, highlighting the motor cortices. The actual motor cortex F1 lies anterior to the central sulcus and is colored orange. All shaded areas anterior to the primary motor cortex are so-called premotor cortices or (pre)supplemental motor areas. These areas are heavily interconnected with area F1. For an interesting overview on how these areas relate action and perception to one another, see [12]. For our purposes, the focus is going to be on one specific premotor area, area F5, which is of particular interest to this work.

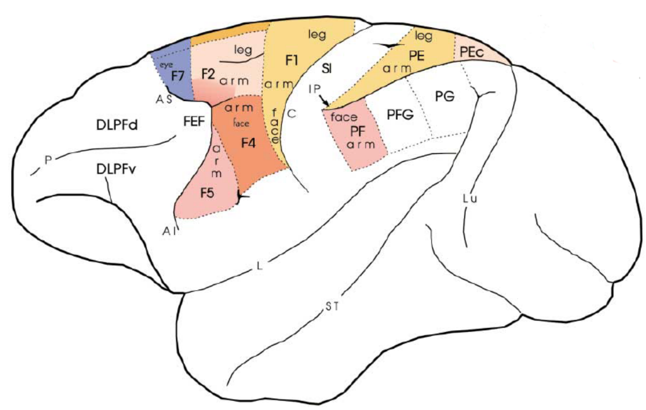

Figure 3: Lateral view of an adult Macaque Monkey brain, adapted from [11]

Area F5 becomes active when the monkey performs certain actions. Specifically, in addition to mouth-related actions, the neurons in area F5 have been found to primarily encode hand-related actions such as holding, tearing or grasping.

Figure 4[3] shows typical discharge patterns of neurons when the monkey performs such actions. Figure 4b and Figure 4a show recordings from two different neural populations found in F5, canonical [11] and mirror neurons [13].

These two groups do not differ with respect to their motor receptive fields: both canonical and mirror neurons can be activated through having the monkey execute different actions. Figure 4a show recordings of a mirror neuron while the monkey grasps a raisin in the dark, thus not receiving any visual stimulation. Similar results are to be expected from canonical neurons.

Figure 4b shows the response of a canonical neuron while the monkey is grasping either a ring, cylinder or sphere. Note that this neuron's responses do not reflect the superordinate category of grasping movements, but the specific type of action required to grasp either of the objects successfully. However, this does not need to be the case: both canonical and mirror neurons with motor receptive fields for movements of different levels of abstraction have been found.

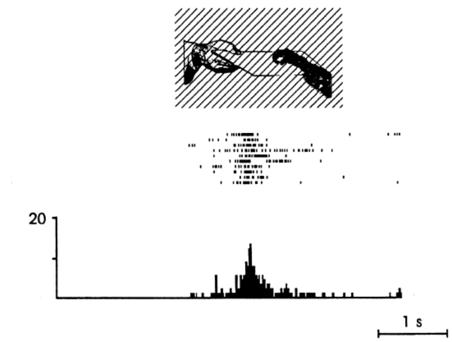

(a) Response of a F5 mirror neuron while the monkey is grasping an object in the dark.

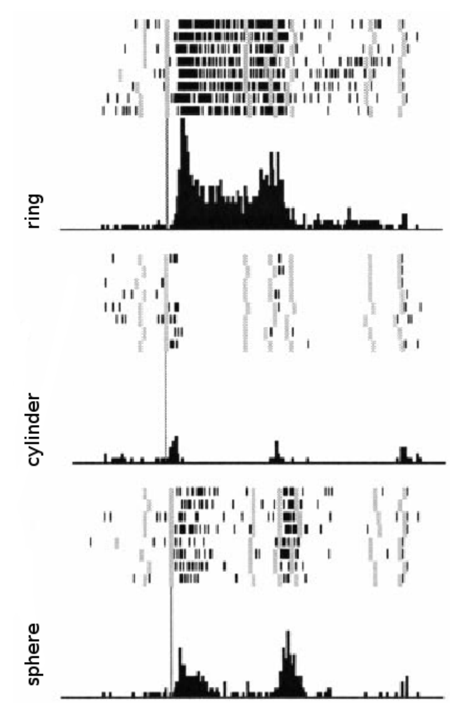

(b) Response of a F5 canonical neuron while the monkey is grasping different objects.

Figure 4: Recordings from a mirror neuron while monkey was grasping in the dark

[3]All spiking pattern figures are adapted from [12].

Though similar in their motor response properties, canonical and mirror neurons differ radically when it comes to what visual stimuli are required to elicit a response: while canonical neurons can be activated by the mere sight of an object suitable for a certain action, mirror neurons only react to the sight of goal-directed actions executed by others. Mirror neurons remain silent when the animal is just exposed to the object such an action is directed towards. These circumstances are illustrated in Figure 5 and Figure 6.

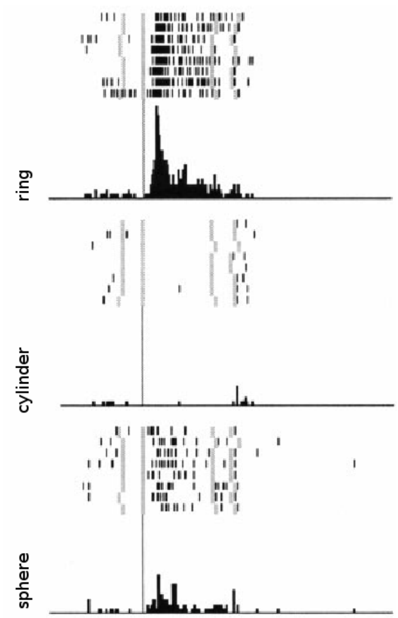

Figure 5: Result of visual stimulation of an F5 canonical neuron

The visual response properties of a canonical neuron for the same set of objects as in Figure 4b are shown in Figure 5. Here, the monkey's task was merely to fixate the objects in question, as opposed to being instructed to act on them. When comparing the responses, we see that the recorded neuron preferred the ring, and slightly less the sphere in both motor and visual conditions. This canonical neuron is thus said to have congruent motor and visual receptive fields.

In contrast, the conditions by which activity of mirror neurons is elicited is shown in Figure 6. Here, we see that the recorded neuron is triggered by the sight of one specific grasp type (Figure 6a), but is not responsive to other types of manipulation geared towards the same object (Figure 6b). By comparing these patterns to the grasping in dark condition of Figure 4a, we see that also this neuron has congruent motor and visual receptive fields.

In summary, mirror neurons and canonical neurons are indistinguishable on the basis of their motor response properties, but differ radically in their visual response properties. Mirror neurons have been connected to the action recognition and action understanding capabilities of primates from very early onwards [14]. On the other hand, the role of canonical neurons has been interpreted in different ways. One of these interpretations, which we will discuss in more detail later, is that canonical neurons function as a filter in order to narrow down what actions could possibly be applied to a given object [15, 2]. Canonical neurons could thus be seen as providing the neural basis for the psychological concept of affordances [12].

(a) Recordings of a single mirror neuron while the monkey is either observing or executing the same action.

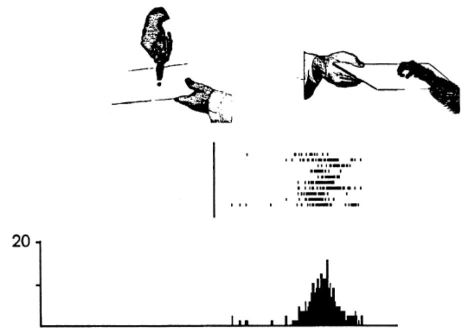

(b) The same mirror neuron is not activated by the sight of a different grasp.

Figure 6: Visual receptive field of a Mirror Neuron

1.3 Goal Statement

This thesis' broader domain is how to enable a robot to choose a good grasp to apply to an object in the light of its features. This is addressed using a probabilistic affordance based approach. Specifically, the work presented in this thesis is based on research by Montesano et al. [1] and Metta et al. [2]. Montesano and coworkers recently introduced the idea of using Bayesian Networks as a formalization tool for affordances. Bayesian Networks were trained to relate actions, objects and effects or outcomes to one another. The action primitives of grasping, tapping and touching an object which were used in their 2008 paper are distinctly different. Since this thesis deals with variations of grasping movements only, our first aim is to replicate the positive results obtained with Bayesian Networks in the domain of grasping. Additionally, a set of more variable objects is used.

Trained Bayesian Networks can be used, among other things, to infer the marginal probability distribution of a certain effect given an object and an action. Albeit useful, this probability distribution cannot be considered in line with the mainstream psychological notion of affordances, which is closer related to the "appropriateness" and "commonness" of an action given an object. This emphasis is more aptly captured by the work of Metta et al., who refer to the marginal probability distribution over actions given an object as affordance. The second goal of this thesis is thus to demonstrate a technique for arriving at the affordance representation introduced by Metta et al., based on the available Bayesian Network.

2 Basic Concepts of Bayesian Networks

Let us now turn to a short review of the main tool used for modeling throughout this work: Bayesian Networks, from now on abbreviated as BNs. BNs are a popular tool in the AI uncertainty community. Charniak [16] lists as representative problems BNs have been applied to medical diagnosis, map learning, language understanding, computer vision and heuristic search. The following section serves to make the reader familiar with the basic concepts of BNs, especially w.r.t. to the steps needed in order to use them as a model for affordances. Most of the information presented in this section is adapted from [16, 17, 18, 19].

2.1 Introduction & Definition

A Bayesian network is defined as a directed acyclic graphs (DAG) $G = (V, E)$ where V denotes vertices or nodes and E denotes edges. The nodes $x_v \in V$ represent random variables (RVs), thus we will use the term "node" and "RV" interchangeably in the following. For each of the random variables in V it must hold that

$$P(s) = \prod_{x_v \in V} P(x_v | x_{pa(x_v)}) \tag{1}$$

where $s = \{x_1, ..., x_n\}$ and $x_{pa(x_v)}$ is the (possibly empty) set of parent nodes of x_v. Thus, the key property of BNs is that the joint probability distribution of all their RVs can be obtained from the individual probability distributions of each RV conditioned on its parents, instead of all other RVs in the BN. This statement is equivalent to saying that each RV is independent of its descendants, given its parents:

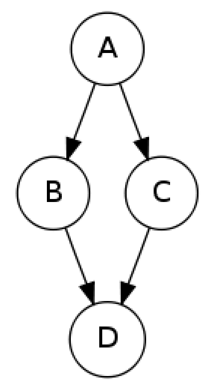

Figure 7: A simple DAG

$$P(x_v, x_{des(x_v)} | x_{pa(x_v)}) = P(x_v | x_{pa(x_v)}) \cdot P(x_{des(x_v)} | x_{pa(x_v)}) \tag{2}$$

where $P_{des(x_v)}$ denotes the (possibly empty) set of descendants of node x_v. This latter formulation emphasizes that BNs fulfill the local Markov Property. The strength of BNs lies therein that they can efficiently represent joint probability distributions by exploiting the statistical independence relationships between their nodes. The statistical independence relationships encoded in a BNs go much beyond the simple and obvious lack of edges between two nodes. More specifically, in order to find out whether two RVs x_1 and x_2 are statistically independent, conditioned on the possibly empty set of evidence E, we need to find out if x_1 and x_2 are d-separated. Although the concept of d-separation is crucial for the theory of BN, we refer the interested reader to [17] for an intuitive explanation.

After having defined what we expect from a BN, let us now consider how to put the advantages of a sparse representation of joint probabilities to practical use. When confronted with a modeling problem in real life, usually all we can hope for is that we are given a complete[4] data set D comprising N observations of all variables $x_v \in V$. From this data set, we have to:

1. identify the structure of the network, which corresponds to learning $x_{pa(x_v)},\ x_{des(x_v)} \quad \forall x_v \in V$.
2. learn the networks parameters, which corresponds to learning $P(x_v|x_{pa(x_v)})$.
3. possibly do inference, i.e. infer the value of one or more RVs, given that we already know some evidence E.

2.2 Structure Learning

The first thing to do when modeling with BNs is to identify the underlying graph structure which presumably gave rise to the observed data set D. In a bayesian setting, we can phrase this more formally as identifying the graph G^* whose probability given the observed dataset D is maximal. Thus we want to find an optimal graph G^* with

$$G^* = arg\ \underset{G}{max}\ P(G|D) = arg\ \underset{G}{max}\ \frac{P(G) \cdot P(D|G)}{P(D)} \tag{3}$$

where Bayes' Theorem was applied in the last step. Unfortunately, the number of possible structures or DAGs is superexponential in the number of nodes [20] [21] and thus an exact enumeration and scoring of all DAGs is not feasible.

Nevertheless, there exist multiple algorithms to either search or sample the space of all possible DAGs efficiently. In the following, we will shortly introduce the reader to the concepts of the algorithms used later in this work. These are K2, a local greedy search algorithm [22], a dynamic programming (DP) approach [23] [24], and DP combined with MCMC sampling [25].

If a topological ordering of the nodes in a BN is known, local search algorithms such as K2 can be effectively applied. K2 starts out with assuming that no node has

[4] Several complications arise if only incomplete data is available, but these issues are disregarded for now.

a parent. It then iterates through the given ordering, and for each node, it adds the parent node that most increases the score[5] of the then obtained graph.

If such an ordering is not known or the goal is to obtain a probability distribution over possible graphs rather than a point estimate such as the one provided by K2, one has to resort to global search strategies. In this domain, the best established methods for BN structure learning are based on Markov Chain Monte Carlo (MCMC) sampling, and, more specifically, the Metropolis-Hastings algorithm. Metropolis-Hastings samples the distribution over graphs conditioned on D. This distribution over graph structures can be of interest if the goal is not to find the single DAG maximizing the likelihood of the data, but only in a set of key structures shared among multiple probable graphs. Furthermore, this distribution can be helpful in cases where it is important to keep a set of alternative structures for future evaluation.

Koivisto and Sood [23] first proposed an algorithm to estimate the exact posterior edge marginals[6] and thus the globally optimal network based on DP techniques. This was further refined in [26]. The trade-off for the exact computation of edge marginals, however, is that we have to make assumptions about the structure of the graph prior $P(G)$, which induces a bias in the obtained result (see [25] for details). When n is the number of nodes in the graph, the DP algorithm has a complexity of $O(n2^n)$, so it is limited in its applicability w.r.t. larger problem domains.

Finally, Eaton and Murphy [25] proposed a hybrid approach between DP and MCMC sampling. By using the edge marginals calculated using DP as a proposition distribution for MCMC sampling, the number of samples to be drawn is drastically reduced. Moreover, the bias introduced through DP methods can be counterbalanced. However, as this procedure relies on DP methods as well, the proposed algorithm does not scale well if there are more than about 20 RVs.

2.3 Parameter Learning

Parameter learning in the context of Bayesian Networks refers to the process of determining the internal conditional probability distribution (CPD) of each node. The form of these probability distributions is not restricted, but theory and inference simplifies if the CPDs of all nodes are multinomial and thus can be represented in tabular form. In addition to the independence relationships encoded by the graph structure these CPDs are necessary for the actual process of statistical inference in the BN.

Given a complete data set, the straight-forward approach to parameter estimation would be to use the maximum likelihood estimate (MLE) of each RV, which in the

[5] One popular choice which is also used in this work is the Bayesian Score.

[6] The probabilities that an edge exists between each two-tuple of nodes (i, j).

case of discrete nodes amounts to simple frequency counts. As a general principle, MLE indeed maximizes the likelihood of the observed data set, but exhibits poor generalization performance. This is partly due to the fact that MLE assumes the probability of all non-observed events to be zero. We can avoid declaring all unseen events to be impossible by using a maximum a posteriori (MAP) estimate instead of the the MLE, introducing suitable priors on parameters. Such MAP estimates smooth the MLE by pretending that each possible case was seen a specific number of times in the data set. The conjugate prior of the multinomial distribution is the Dirichlet distribution. A special Dirichlet distribution is the BDeu[7][27] which ensures that the likelihood equivalence principle[8] is fulfilled.

2.4 Statistical Inference

Inference in a Bayesian network refers to the following task: some of the values of the RVs in a network are already observed, and we are interested in the values of some or of all of the remaining RVs.

Let us illustrate why this process is not as trivial as it seems to be in such a plain network as the one shown in Figure 2.1: suppose we observed the value of the node D and wanted to know what impact this has on the conditional probability distribution at node C. Knowing D has a twofold influence on C: via the direct connection from C to D, but also indirectly, because knowing the outcome of D influences our knowledge about node A which in turn projects to C.

In principle, the problem of inference in Bayesian Networks is NP-hard. However, for a restricted class of graphs, inference can be done in linear time with respect to the number of nodes. This class of algorithms operates on singly connected graphs, so named because for all pairs of nodes $u,\ v \in V \times V$ there is at most one simple path[9] from u to v. The key to most exact inference algorithms is thus to transform the given network into a singly connected graph first. This is often achieved by using clustering, where nodes in the network are combined or merged into one until the resulting graph is singly connected. If the resulting graph moreover fulfills some additional constraints, exact inference is feasible. A well established algorithm of this class is the so-called Junction Tree Algorithm (JTA) [28] [29], which we will use throughout this work. JTA works by transforming the BN into a junction tree and subsequently performing exact inference using belief propagation.

Alternatively, there are a lot of exact and approximate algorithms available for

[7]Bayesian Dirichlet uniform equivalent

[8]Equal likelihood of models in the same Markov Class

[9]A path were no vertex is visited more than once.

inference in BNs, some of which can be used for inference in real time systems[30].

2.5 Importance of Interventional Data

Under certain circumstances, BNs can be used for modeling causal relationships [18]. Then, a directed edge from vertex X to vertex Y is interpreted as X causing Y. This, however, does not need to be the case. As an example, consider Figure 2.5, which shows three graphs over the variables X, Y and Z: they all encode the same independence conditions, namely that X is independent of Z given Y. Thus, all DAGs in 2.5 are said to be Markov equivalent and cannot be distinguished on the basis of observational data only. Observational data is on hand if the outcomes of each random variable is only observed, but not actively intervened upon. By contrast, data is said to be interventional if certain nodes were forced to take on specific values. Accordingly, when experiments are performed, and thus interventional data can be used, it is possible to make statements about the causal relationships between variables in a BN.

All structure learning algorithms introduced earlier can make use of interventional data. This amounts to an algorithm having additional information about which outcomes were intervened upon. When a node is intervened upon, it can be thought of as effectively cut off from its parents. This yields independence information to the structure learning algorithm.

A different type of a priori knowledge concerns the temporal sequence of observations of RV outcomes. When it can be assumed that the RVs in a BN are causally related and temporal information is available, this knowledge can be incorporated via the process of layering. Layering aims at modeling the temporal sequences of RVs by grouping all RVs in a given BN into layers and restricted the number of connections between layers to forward connections, i.e. connections from lower-numbered to higher-numbered layers. Thus, layering can be considered as amounting to the assumption that there is no such thing as causation going backwards in time.

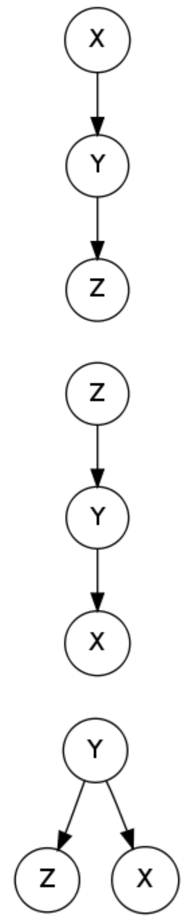

Figure 8: All of the above DAGs encode that $X \perp Z|Y$

3 Related Work

3.1 Affordances as Relations: Bayesian Networks

Indispensable for this thesis is the work of Montesano, Lopez and coworkers [1, 31]. In their recent papers, they explicitly addressed the issue of affordance representation and quantification. To the author's knowledge, they were also first in using Bayesian Networks for affordance modeling.

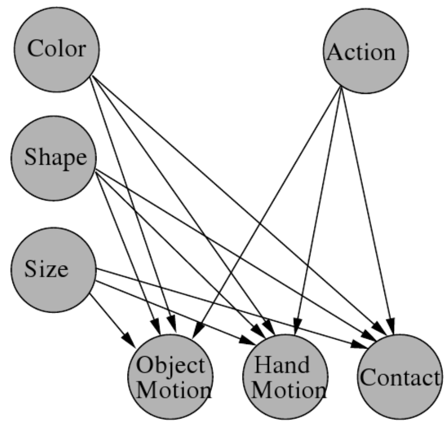

(a) Generic Network

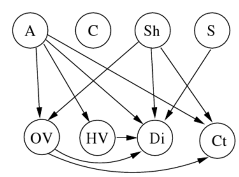

(b) Network obtained with K2

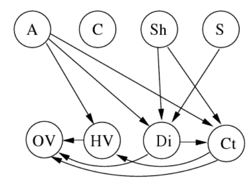

(c) One of the networks obtained with MCMC

Figure 9: Networks by Montesano et al. In 9b and 9c, the node Di encodes direction of movement of the object.

In a series of experiments on affordances, world interaction and imitation learning, Montesano et al. had their robotic platform BALTAZAR apply several actions[10] to a set of different objects and observe the effects or outcomes[11] of these actions. Each of the object features, actions and perceptual effects was assigned a RV and subsequently modeled as a node in a Bayesian Network. They exemplified the capabilities of the learned networks in the domains of for effect prediction and imitation learning. The generic layout of all BNs used in their work is illustrated in Figure 9a, showing the object features $F_1, ..., F_N$ to the left, the effect RVs $E_1, ..., E_N$ at the bottom and the action in the upper right-hand corner. Typical examples of trained networks are shown in Figures 9b and 9c.

Interestingly, Montesano et al. could generalize their idea of using BNs for affordance learning learning the free parameters of motor primitives themselves [1]. In the case of a grasp such free parameters correspond to the angles of each joint that are controlled by the robot itself. It was shown that by modeling certain free parameters such as the height of the arm as a RV itself, the correct height for a given object could

[10] These were grasping, tapping and touching

[11] For example duration of contact with the object and the relative hand-object velocity

be inferred. Although this thesis treats different grasps as motor primitives, it is in principle possible to model actions as decomposable w.r.t. certain features. Then, motor features such as the height of the wrist at the moment of grasping, could serve to improve the flexibility of the learning system.

3.2 Probabilistic Framework for Mirror Neural Activity

A different approach to affordance formalization in the robotics domain was taken by Metta and colleagues. In a paper expressively titled "Understanding mirror neurons: a bio-robotic approach" [2], they proposed a general probabilistic framework for modeling the functionality of the premotor area F5, where mirror neurons as well as canonical neurons reside in the brain of adult macaque monkeys. It was suggested that this system - which is inherently adapted to the generation rather than recognition of actions - can be readily coopted for action recognition. This idea is based on neurophysiological studies, such as the work by Fadiga and Rizzolatti reviewed earlier.

Here, the activity of the mirror neuron system is represented as the probability of an action a_i out of a finite motor repertoire A, given certain motor features/parameters f of an action and information about an object o_k drawn from a finite set of objects O. The likelihood of the activation of the mirror neuron system, can then - using Bayes' proportional law - be expressed as

$$P(a_i|f, o_k) \propto P(f|a_i, o_k) \cdot P(a_i|o_k)$$

where

- $a_i \in A$ is a specific action out of a given set of actions
- f are the motor features specifying a_i
- $o_k \in O$ is a specific object out of a given set of objects

The response of the canonical neurons in F5 - and thus the affordance coding this thesis is interested in - is thought to correspond to the term $P(a_i|o_k)$, i.e. the likelihood of invoking a certain action given an object. In short: if presented with a sphere, one is much more likely to grasp it with a precision grip whereas the handle of a hammer is much more suited for a power grasp. The extraction of affordances is assumed to take place "automatically", which conforms to neurophysiological experiments showing that canonical neurons fire when presented with an object that would trigger a certain action, although the animal was not instructed to act upon this object.

The conditional probability of an action given an object is an intuitive, precise and practical formalization of the concept of affordances. Furthermore, it is not assumed that affordances are unambiguous, since certain object features, such as a small elongated box, might yield high probabilities for e.g. both precision and power grasp. Furthermore, representing affordances as probability distributions eases their use in robot control.

In their experimental work Metta and coworkers focused on developmental approaches [2, 32] in order to elucidate how the representations found in premotor cortex could come about without imposing implausible constraints on learning mechanisms. However, to the author's knowledge, no systematic attempt at "filling in the prior" $P(A|o)$ has been made.

4 Data & Methods

4.1 Experimental Paradigm & Data Acquisition

All data was collected using the humanoid robot iCub [33] shown in Figure 10, an open-source robot developed in the EU Research Project RobotCub with the specific aim of providing a platform for experiments in cognitively developmental robotics. iCub was designed to resemble a three year old human child, with similar physique and degrees of freedom. Each of iCub's arms has 16 degrees of freedom, seven of which are assigned to the portion from shoulder to wrist and nine to the control of the hand. The experimental protocol is outlined in Table 1:

stage	activity
1	bringing the hand to the object
2	grasping the object with one of four different grasps
3	probing the object stability with one of two different moves
4	assessing the grasp stability
5	releasing the object
6	bringing the hand back to the starting position

Table 1: Experimental Protocol

Four actions, all considered primitives, were designed. In order to be better able to evaluate the affordance prediction capabilities of the resulting system, all four actions were preprogrammed and furthermore tailored to one specific reference object for which they should yield maximal stability. Table 2 summarizes all four grasps along with their reference objects:

grasp	encoding	reference object
large power grasp	1	spindle
small power grasp	2	cube
large precision grasp	3	softball
small precision grasp	4	textmarker

Table 2: Overview of Grasp Types

All grasps were categorized according to Napier [34]: the term precision grip refers to grasps where the thumb and the finger tips make up the opposition space whereas

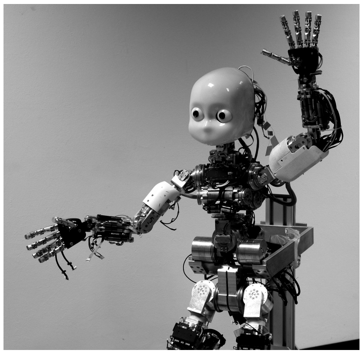

Figure 10: The humanoid robot iCub

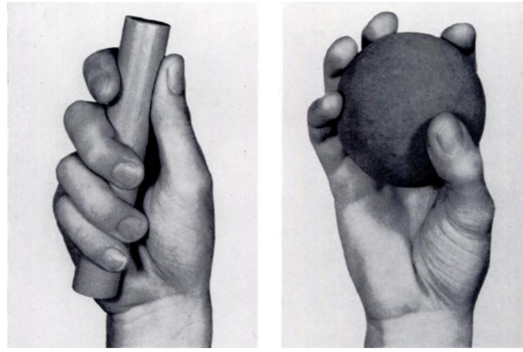

Figure 11: Power and Precision Grasp, adapted from Napier [34]

power grasps are characterized by the opposition of the fingers to the palm. This difference is illustrated in Figure 11. Furthermore, for each opposition type, small and large grasps were distinguished based on the degree of closure of the proximal joints of the fingers when the grasp was fully executed. The distinct feature of opposition space could be easily controlled in the process of designing the primitive actions by carefully setting the opposition angle of the thumb. Joint angles for all grasps are given in Table 15 on page 51. In addition to the grasp type, the second component making up an action that iCub applied to each object was a probing movement used to assess the stability of a given grasp/object combination. Two probing actions were implemented:

probe	encoding
rotation of the forarm of approximately 150 degrees	Rot, 1
up and down movement of the shoulder and elbow joint	UpD, 2

Table 3: Probing Movements

As mentioned, the probing movement served to assess the stability of the grasp. Naturally, there are different possibilities for defining grasp stability. The most prominent measures for grasp stability are force and form closure [35] and are drawn from an engineering perspective. Force closure and the related form closure are calculated by analyzing the forces exerted from the finger tips or the palm assuming certain friction coefficients and defining a stable grasp as one where arbitrary perturbations can not alter the position of the grasped object. These measures, however, are unsuitable for our purposes. First, the geometric analysis and knowledge about friction coefficients are not an approach which would be considered valid in a developmental setting. Instead, it is desirable that the robot learns what stability refers to from its own experience, similar to a child who learns to associate certain proprioceptions with consequences on the external world, such as the pressure exerted on the palm while steadily holding an object. Unfortunately, the tactile sensors needed for this kind of proprioception were not available on the robotic platform iCub. Thus, we could not make use of information provided e.g. by pressure sensors to infer the grasp stability. Instead, this information was provided from the outside, i.e. the experimenter categorized each performed grasp as either stable, unstable, or failed[12].

Finally, an overview of the 21 objects used throughout the experiment is given in Table 16 on page 52 and Figure 19 on page 53. The objects were selected to represent a wide range of features, but only objects which were in principle - with one grasp or another - suitable to be held by iCub's hand were used for the final data collection. Each action, i.e. probe and grasp combination, was applied to each object between three and ten times. A total of 758 trials was collected.

4.2 Data Preprocessing

In principle, the nodes or RVs of BNs can hold probability distributions of any form. Since BN theory is best studied for multinomial RVs and also the neurophysiological evidence reviewed in Section 1.2 hints to some discretization of perceptual input, all information was discretized before training the BNs. For grasp types and probing movements this was not needed, since they were already designed to be multinomial RVs. Also the stability RV, which was judged by an external observer, was already a multinomial RV. Thus, only the RVs corresponding to object features remained to be discretized. The chosen features were object shape, height, color, texture and rigidity, each of which could take on multiple values. Please see 13 on 51 for details on these object features. Montesano and coworkers [1] directly applied clustering algorithms to discretize perceptual information. For the purpose of this work, each object was manually assigned the values of all features. A summary of all objects and their feature values is given in Table 16 on page 52[13].

After preprocessing, the data was available in a format where each of the trials corresponded to an eight tuple made up of five values for object features, two values

[12]However, it could be shown that tactile and proprioceptive information is very useful in object categorization when using a SOFM [36], and thus it may be possible that subsequent work succeeds at replacing the external categorization of grasps by tactile and proprioceptive data.

[13]The choice of object features is an important and crucial problem in its own right, but beyond the scope of this work. There exist many promising computer vision algorithms which could aid this purpose, see e.g. [37] for a method suitable for application in robotics.

for action encoding (grasp type and probing movement), and one for the observed stability outcome. Note that in the following the term action refers to a specific grasp and probe combination, deviating from the its everyday use.

4.3 Learning the BN Topology

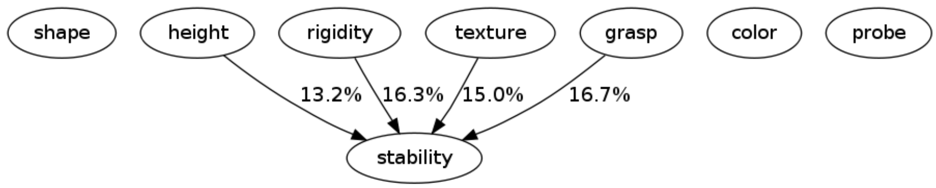

(a) Topology uncovered by DP and DPMCMC

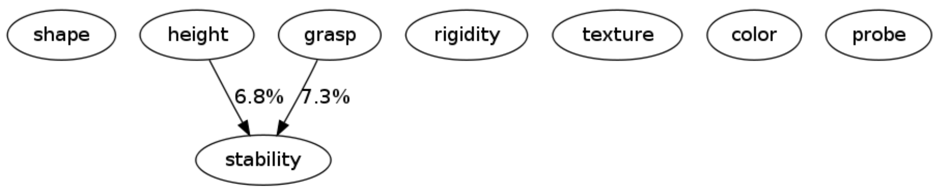

(b) Topology uncovered by K2

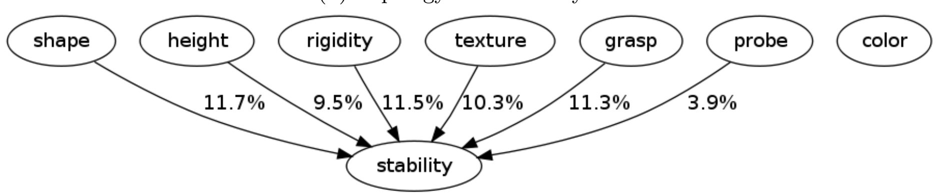

(c) Manually designed topology

Figure 12: Different BN topologies; edges are annotated with true average link strength, a concept explained in Section 4.4.

Several methods for structure learning in BNs were examined. First, pure Dynamic Programming (DP) [26] and a hybrid approach using DP to restrict the search space which is subsequently sampled using Markov Chain Monte Carlo methods (DPMCMC) [25] were examined. Both of these approaches are subsumed under the notion of “DP methods” in the following. All algorithms here made use of interventional data, assuming perfect interventions. DP methods are evaluated in Section 4.3.1. Section 4.3.2 examines the greedy K2 algorithm. Finally, Section 4.3.3 illustrates a manually designed topology. Section 4.3.4 illustrates the point that interventional data is needed.

4.3.1 DP methods

Both DP and DPMCMC can compute the marginal probabilities of a directed edge between all ordered pairs of RVs. For the MCMC step of the hybrid approach, 5000 samples were drawn with a burnin period of 500. Figure 13a on page 21 shows the edge marginals obtained with DP; those of DPMCMC are similar. The heat map illustrates that the DP methods correctly identified grasp stability to depend on the object features height, rigidity, texture and furthermore on the specific grasp type used. Surprisingly, the probing action is found to have no influence on stability. The same holds for the shape and color of the object. Due to the unbalanced number of trials performed with each object/action combination, several edges are indicated with an approximate probability of 0.3.

The same overall structure is obtained when layering constraints are used in addition to interventional data. Layer I was assigned to hold all object features, layer II all action features and layer III held the only effect of stability. No connections within or between layers I and II were allowed; the number of edges projecting from layer I or layer II to layer III was not restricted. The connectivity pattern did not change when layering was used. Instead, the uncertainty about the existence of edges between and within object and action features is removed. The BN topology obtained through DP methods is illustrated in Figure 12a.

4.3.2 K2

Structure learning using the K2 algorithm [22] was performed. K2 is a local, greedy search algorithm in the space of DAGs. K2 needs to be initialized with an ordering of the RVs. It works by incrementally adding parents to each node if this increases the score of the resulting structure. The node ordering used was shape, height, rigidity, texture, color, probe, grasp, stability. It was found that the network topology found by K2 is much sparser than the one obtained through DP methods, having only two edges: one from grasp type to stability and one from object height to stability. This topology is shown in Figure 12b

4.3.3 Manual Design

Given that the number of RVs for the structure learning task at hand is quite small, we can tell intuitively

1. that some object as well as action features should have an impact on stability.
2. that action features and object features are in general independent of each other.

Using these two constraints, it is possible to design a plausible BN topology with dependencies of grasp stability on all other object and action features except for color. We will keep this topology, shown in Figure 12c, for later comparison purposes. In the following the BN obtained from this topology will be referred to as “EN” (engineered) BN.

4.3.4 Interventional Data

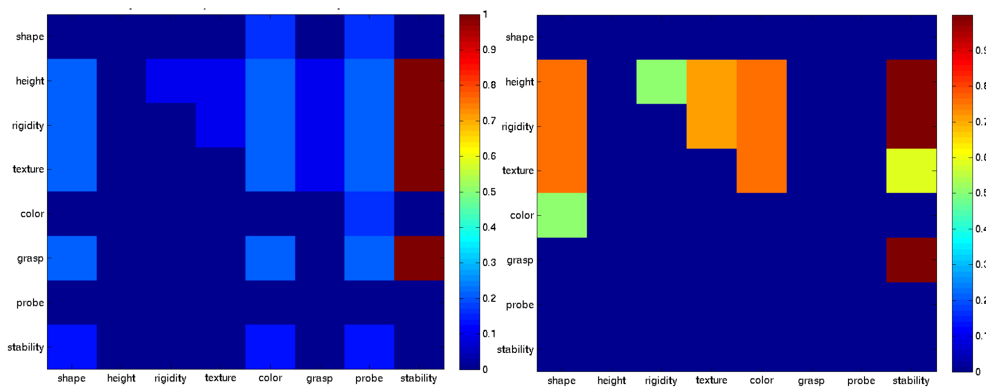

(a) Computed with Interventional Data (b) Only Observational Data used

Figure 13: Edge marginals obtained with DP. Color indicates the probability associated with a directed edge from the RV on the y-axis to the RV on the x-Axis. The order of the RVs from top to bottom and left to right is shape, height, rigidity, texture, color, probe, grasp and stability.

Except for the EN network, all computations of edge marginals mentioned above made use of interventional data or layering information. In order to illustrate the importance of either type of information in recovering the structure of the network which presumably generated the observed data, Figure 13b shows the edge marginals as estimated by DP when all variables are treated as observational. Note that the probabilities of edges connecting object and action features are relatively high when compared to Figure 13a. In order to be able to interpret the resulting BN as encoding causal relationships, it is thus crucial to either gather interventional data or to take into account the temporal sequence of events through layering.

4.4 Parameter Learning

The internal parameters of the nodes were approximated through MAP estimation using BDeu priors initialized with a weight of 10 percent of the whole data set, which amounts to 7.5 pretended observations for each unseen event.

The edges of the graphs in Figure 12 are annotated with true average link strength information as obtained by the procedure described by Ebert-Uphoff [38, 39]. The true average link strength percentage measure between a node X and Y can be interpreted as indicating how much the uncertainty of the target node Y is reduced by knowing the state of X, if the values of all other parents of Y are assumed to be known.

As we see in Figures 12a, 12b and 12c, all three networks cannot fully reduce the uncertainty about the effect given observed values of all other attributes. The DP network can reduce the uncertainty of the stability node by about 61.20% and the engineered network by 57.49%. The performance of the K2 network is particularly poor with only 14.10%.

4.5 Intermediate Discussion

The BN topologies as computed by DP methods correspond to my subjective experience in that for the specific robotic platform used, stability depends much more on texture and rigidity than for example on the type of probing action. Although the probing actions were intended to exert different forces on the object, the differences between them seem to not have been significant enough to discourage statistical independence between probe and stability. Also, the fact that texture and rigidity are more predictive of the observed stability than the shape of the object (see link strength annotation in Figure 12a), aligns with my observations: very textured and soft objects could easily be grasped by iCub's hand, whereas objects of the same size were commonly failed to grasp when they were too rigid or too smooth. In particular, this is probably due to the material of iCub's hands, which have a smooth surface which does not provide much frictional resistance, unlike for example the human skin. Thus, it is important to remind ourselves that the set of dependencies detected here is to be seen not necessarily as what the experience of a human being would look like, but rather as a reflection of the specific action capabilities and physique of iCub.

5 Predicting Stability

5.1 Evaluation Procedure

As it is often the case with BNs, there is no ground truth on which to evaluate the inference capabilities of any learned network. Yet, it is possible to judge the relative performances of different network topologies using cross validation. In our particular case, we are interested in how good the predicted stability matches the observed stability for any of the computed networks. Note that it would be possible to infer the probability distribution over action features from known object features and stability outcomes as well as any other combination of RVs. However, as prediction capabilities of the networks form the basis for the later computation of affordances, we are going to limit ourselves to the evaluation of the stability inference problem.

To do this, leave one out cross validation was used. The parameters of the three graph topologies shown in Figure 12 were trained on the data from all experimental trials except those trials involving the object used for eventual inference evaluation. Then, each network was queried for a probability distribution Q over stability outcomes for each object/action pairing. From the data spared from parameter estimation, the MLE estimate of the distribution over stability outcomes P was computed. After that, the two distributions P and Q were compared using the KL divergence[14] defined as

$$D_{KL}(P||Q) = \sum_i \; P(i) \; log \; \cdot \; \frac{P(i)}{Q(i)}$$

. Since the KL divergence is not symmetric, note that we consider the inferred distribution Q an approximation to the MLE distribution P.

[14]The KL divergence is minimal, i.e. equal to zero, if P and Q are identical. However, the KL divergence as defined here does not have an upper bound. Furthermore, it's important to keep in mind that the KL divergence is not a metric as it does not fulfill the triangle inequality.

5.2 Results: Relative Performance of different Topologies

Object	S	H	R	T	C	DP	K2	EN
bottle	2	2	2	1	2	10.83	6.36	4.63
brush	1	1	1	2	1	1.92	4.48	3.10
bscarf	4	2	3	3	2	1.59	2.39	2.43
bubbles	3	1	1	1	3	4.37	4.63	4.37
chick	4	2	3	3	1	2.51	2.85	3.56
cube	3	1	2	1	1	5.13	9.02	5.13
cup	4	2	2	1	2	11.77	8.59	4.73
drumsticks	1	1	1	2	1	1.92	4.48	3.10
dvd	2	1	2	1	3	8.09	6.20	8.09
flummi	3	2	1	3	2	4.73	3.12	4.73
marker	3	1	1	1	3	4.37	4.63	4.37
pdice	3	2	2	3	2	3.43	4.90	5.15
pump	2	2	2	2	2	3.20	8.95	6.40
rbubbles	3	1	1	3	1	1.89	5.61	2.35
rmarker	3	1	1	3	3	2.15	6.23	1.81
rubber	3	1	3	3	2	4.80	10.16	4.80
scarf	4	2	3	3	3	1.48	3.83	2.00
softball	4	2	3	2	1	6.67	6.67	6.67
specbox	4	2	1	2	3	6.64	5.07	4.45
spindle	2	2	1	2	1	5.28	6.89	5.35
taperoll	3	1	1	2	1	3.63	4.85	5.36

Table 4: KL divergences between predicted and actual distributions for the DP, K2 and EN networks. The first five columns show the assigned feature values.

Figures 14, 15 and 16 on pages 26, 27 and 28 depict the inferred and MLE distributions for the topologies attained through DP, K2 and manual design respectively. In each subfigure, the left heat map corresponds to the MLE reference distribution P, to which we compare the right heat map showing the inferred distribution Q.

The most striking feature about these figures is that, occasionally, all three networks produce uniform predictions over stability outcomes. In numbers, the K2 network resorts to uniform predictions for 4.7%, the DP network for 33.3%, and the engineered network for 76.2% of all objects.

Table 4 summarizes the KL divergences obtained for all objects when tested with the respective network topologies. Columns one to five repeat the object feature values for easier comparison.

A Friedman's test[15] was performed in order to quantify the differences between generalization performances. The Friedman's test was conducted between the KL divergences obtained by all three network topologies and indicated a significant difference ($P < 0.01457$). Post hoc analysis of the medians for each topology indicated that it is most likely to assume that the DP BN (Median 4.37) performed best, followed by the EN network (Median 4.63), and finally the K2 BN (Median 5.08).

Table 5 shows the ranking of network topologies obtained through simple counting of the number of objects for which the respective networks yielded the minimal KL divergence compared to its competitors. The results are in line with the results obtained through analysis of the KL divergences in that the DP BN performed best, followed by the EN BN and subsequently the K2 BN.

Topology	Smaller	Equal
DP	10	5
EN	4	5
K2	2	1

Table 5: Ranking of BN topologies. The "Smaller" column shows the number of objects for which the KL divergences attained by the respective network topologies were strictly smaller than those of the two competing networks. The "Equal" column lists how often the respective network shared these minima with at least one other network.

[15]The Friedman's test is a nonparametric equivalent of a classical ANOVA. We cannot use an ANOVA in this case since the KL divergence is not a metric and thus our data is on the ordinal level of measurement.

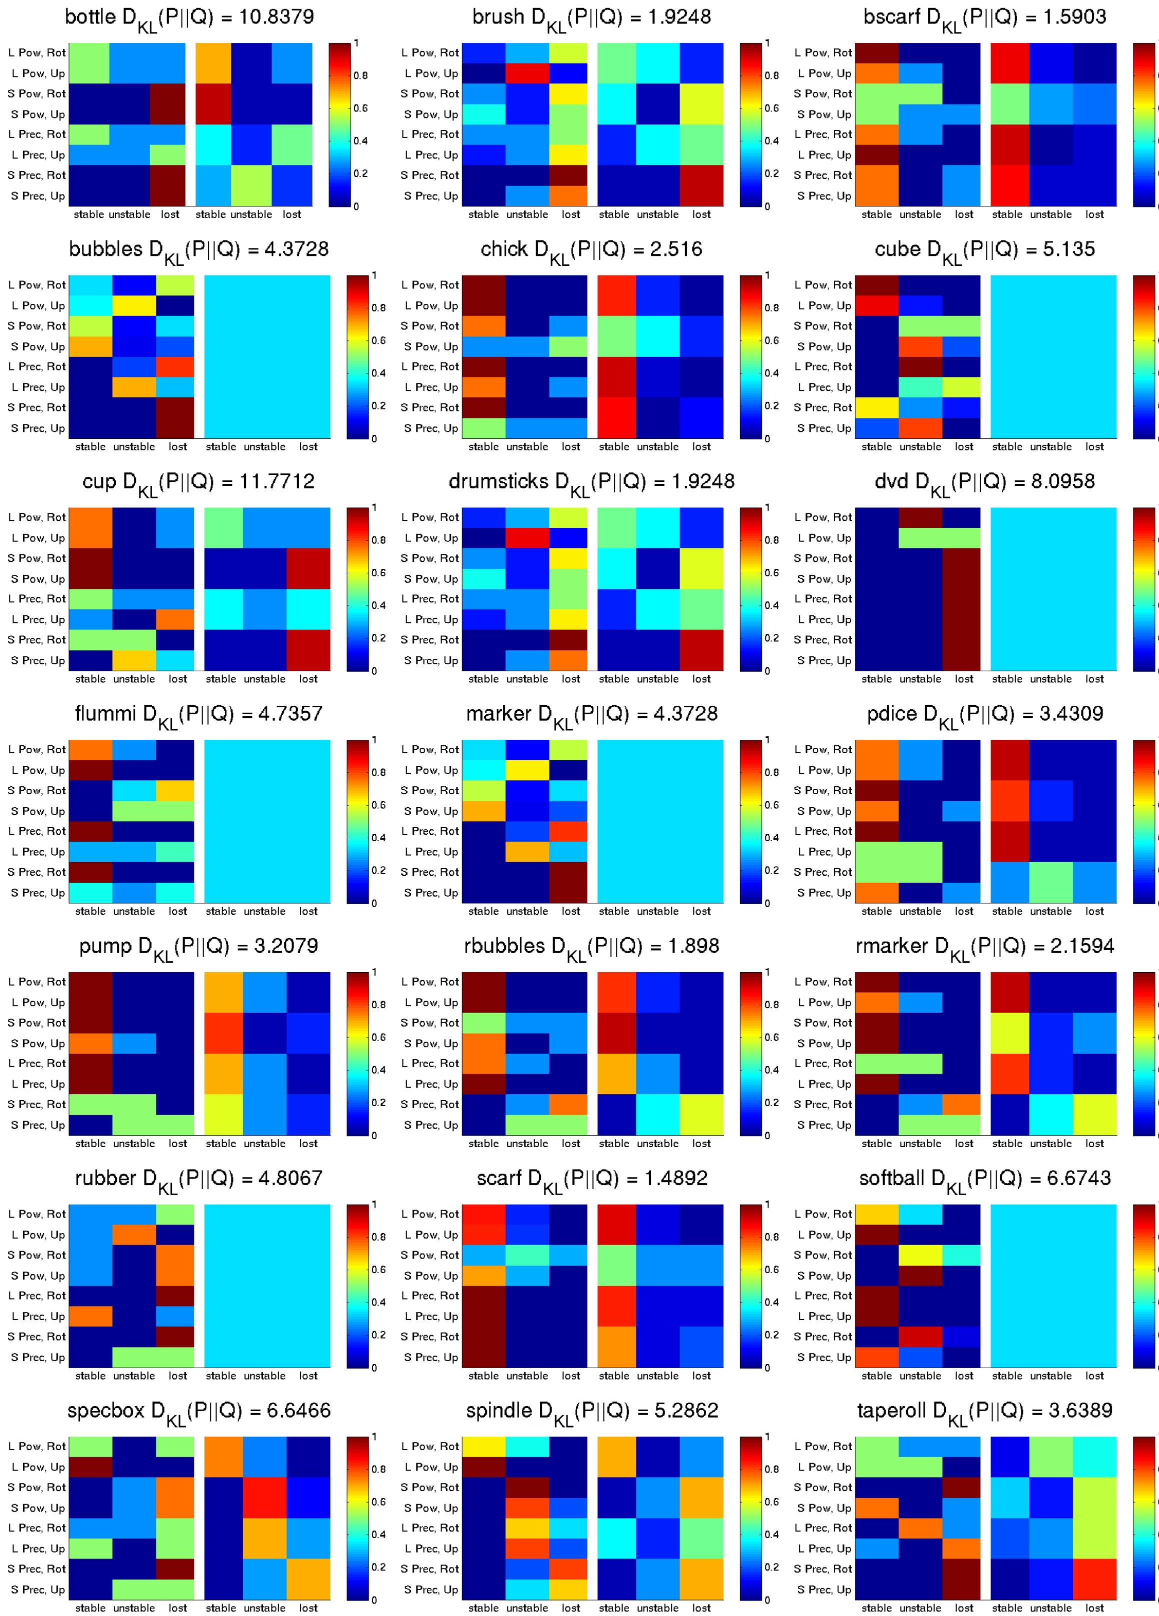

Figure 14: Inference in the DP BN

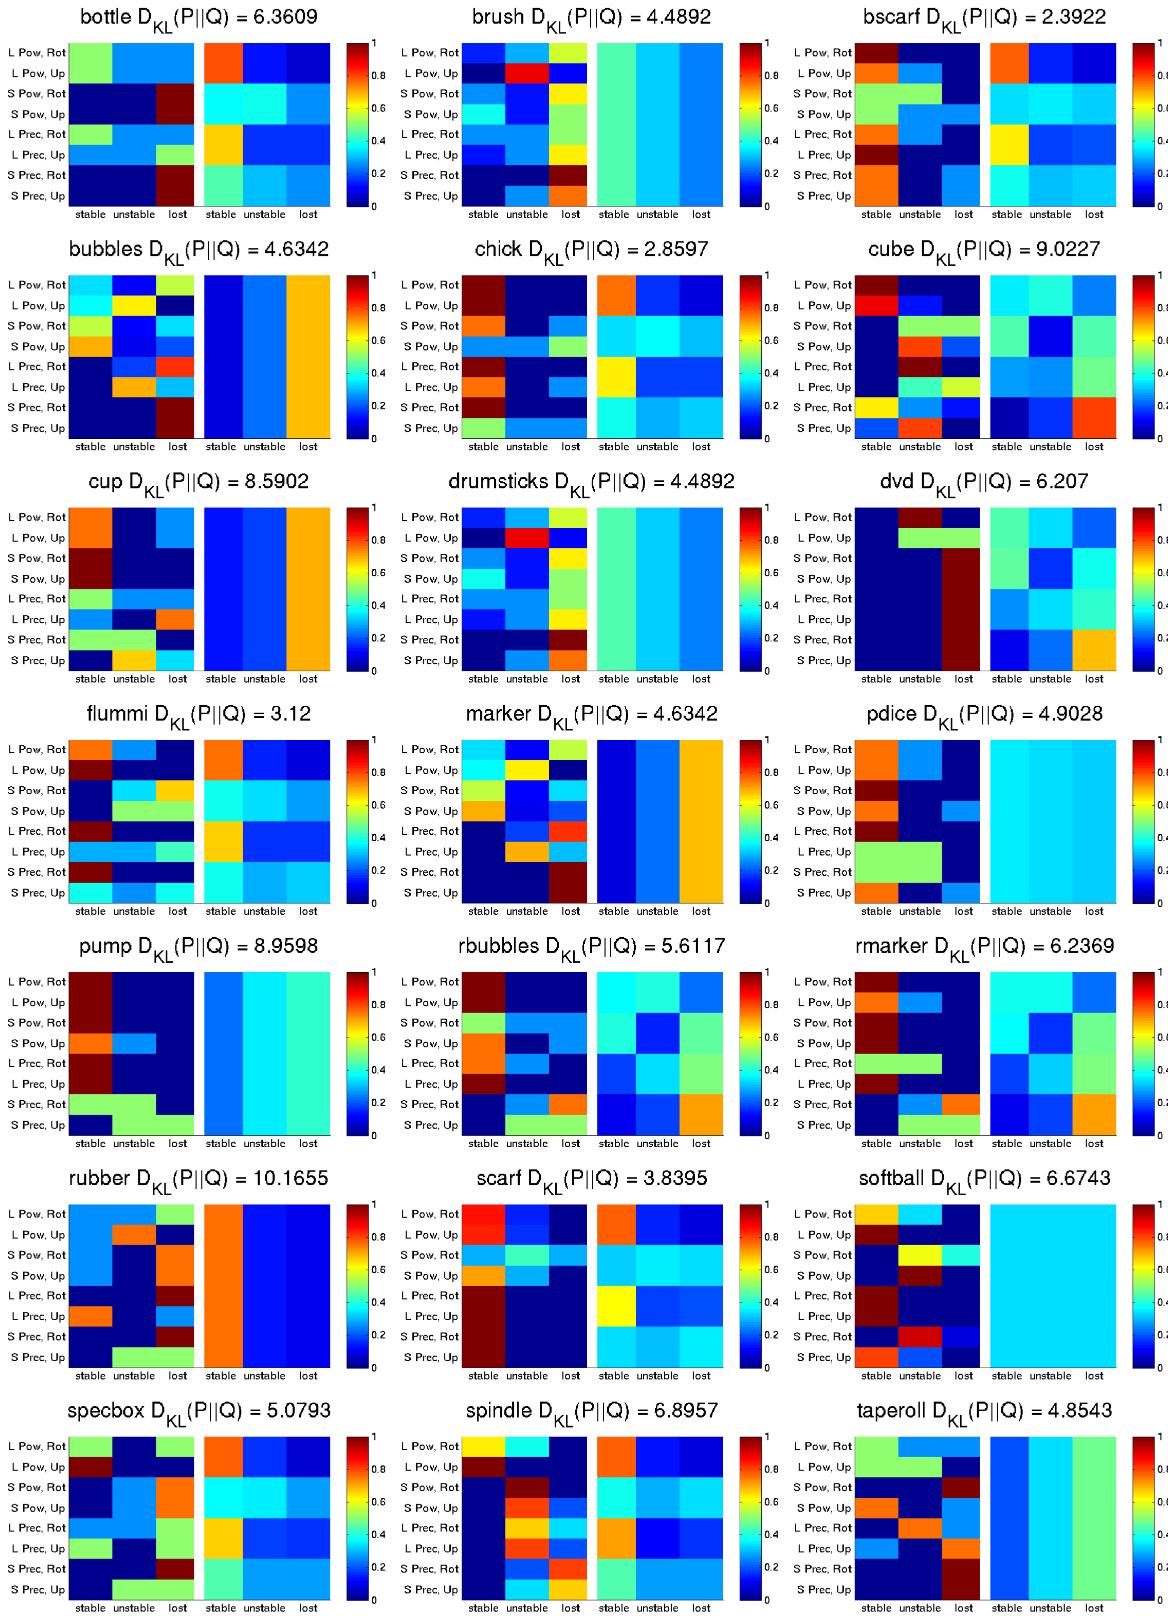

Figure 15: Inference in the K2 BN

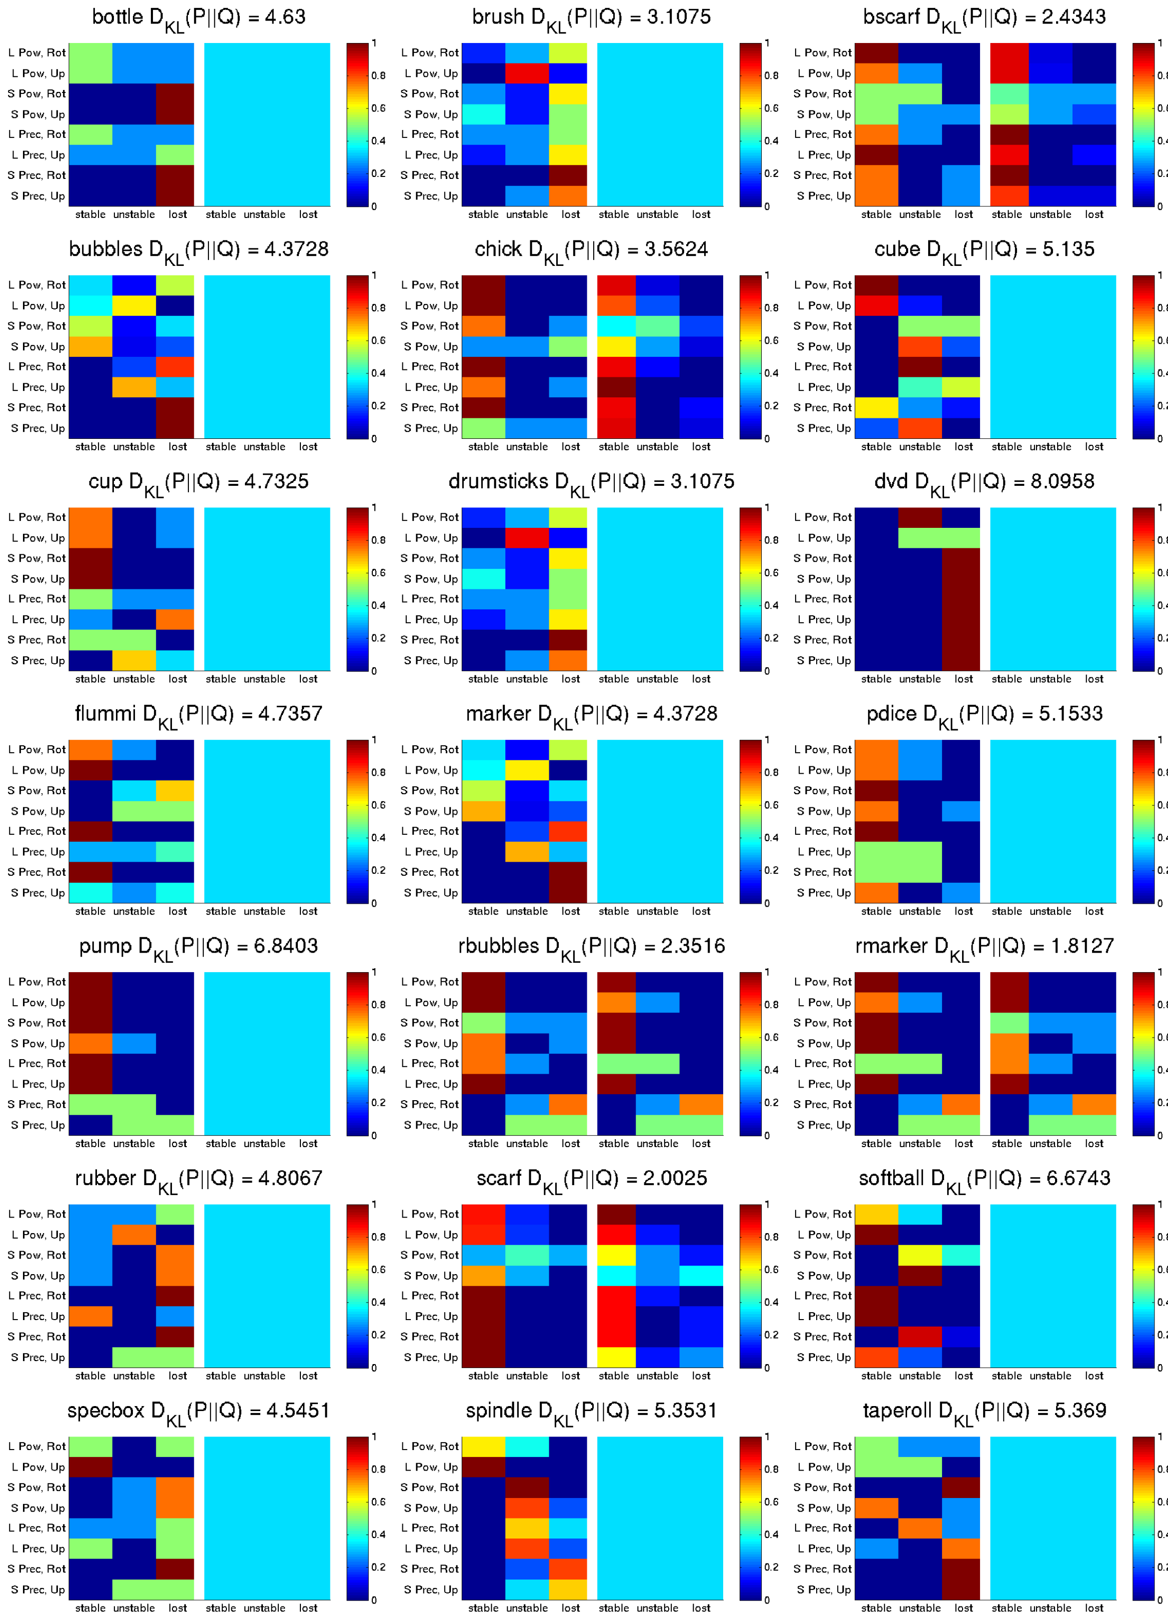

Figure 16: Inference in the manually designed BN

5.3 Discussion

We will begin this section with taking a closer look at the object dvd. First, note that the dvd is a distinguished object in that it was the case only for the large precision grip that the resulting configuration could be called unstable, instead of merely observing a failure to grasp. This can be easily seen from the MLE distribution in Figure 14, page 26. In the particular case of the dvd, it seems that the features used by the EN network and the DP network (both $D_{KL} = 8.09$) were much too specific to accommodate this rather unusual object. The K2 network ($D_{KL} = 6.20$) performed best for this object, although the KL Divergence attained was still relatively high. Only for one further object, the flummi, see Table 4, the K2 BN yielded the best result. This suggests that very unusual and atypical objects seem to be better accommodated by sparse networks, although the evidence can not be considered conclusive[16].

Contrary to the atypical dvd, generalization performance for the object scarf was expected to be fair, since all networks were trained on a very similar object, a blue scarf (bscarf) with identical features except for color. We see that K2, with its "one-dimensional" object concept, fails to exploit this redundancy ($D_{KL} = 3.83$). The EN network ($D_{KL} = 2.00$) performs better but still worse than the DP network ($D_{KL} = 1.48$). We could thus stereotype the scarf as a highly ordinary object whose features are very predictive of its behavior. In this case, the medium number of arcs of the DP network seems to be superior to both high (EN) and low (K2) connectivity patterns.

From these initial considerations it is just a stone's throw to a possible explanation for the percentage of uniform predictions of stability outcomes by the different BNs. Intuitively, these uniform distributions inform us that the respective BN has never been exposed to an object similar to the one which it is queried about. Note that naturally, the BN could not have been exposed to the very same object during training, which is the intended consequence of the test procedure. However, what exactly constitutes a "similar" object, and thus, whether the network's generalization performance is fair, depends heavily on the number and type of features constituting an object for a given BN. Let us recall which object and action features influence stability for each of the three topologies shown in Figure 12 on page 19:

- K2: height, grasp
- DP: height, rigidity, texture, grasp
- EN: shape, height, rigidity, texture, grasp, probe

[16]This is especially true due to the bias of the K2 network to predict nonuniform distributions. This increases the chances that the mentioned scores were merely a product of coincidence.

Proportionally to the number of variables the RV stability is conditioned on, the K2 network resorts to uniform predictions for only 4.7% of all objects, the DP network for 33.3%, and the engineered network for 76.2%. Combined with the above case illustrations, this ordering supports the supposition that the available data is overfitted by the EN network and underfitted by the K2 network. By contrast, the medium connectivity of the DP seems to be suited best for the prediction of stability.

6 Quantifying Affordances?

6.1 Introduction

The preceding section demonstrated the capabilities of the learned networks to solve the task of effect prediction, i.e. the prediction of the stability outcome of a specific grasp/object combination. The performance for effect prediction proved to depend heavily on the network topology and the range of experience available to the learning system. The EN network as well as the network obtained through DP methods, proved to be relatively good at this task, whereas the performance of the K2 network was quite poor. The K2 network is thus dismissed from the following comparisons. Starting from the firm ground the DP and EN networks provided, we are now going to show how these two networks can be utilized to arrive at an alternative probabilistic affordance formalization as presented by Metta et al. [2]. The basis of the methods presented later is that for each object and grasp combination, it is possible to infer the probabilities given in Table 6:

$P(S = stable \mid A = a, O = o)$	probability that the action/object combination will be stable
$P(S = unstable \mid A = a, O = o)$	probability that the action/object combination will be unstable
$P(S = lost \mid A = a, O = o)$	probability that the action/object combination will result in loss of the object

Table 6: Overview of possible stability outcomes

From this CPD, the goal is to arrive at a scoring function measuring the match between each action and object, i.e. the affordance of the object formalized as $P(a|o)$. As mentioned in Section 3.2, the CPD $P(a|o)$ is not only important in its own right, but also integrates into models of the mirror neuron system.

Naturally, how we convert the stability inferences in Table 6 to an estimate of the affordance $P(a|o)$ depends on whether we require a point estimate of affordances (PEA), which would simply amount to the most apt action, or a completely specified probability distribution over actions (PDA). Both results can be important in different contexts. First, we are going to focus on deriving a PDA from the CPD over effects. Afterwards, we are turning towards point estimates, or PEA. Finally, the respective performances of PDA and PEA are evaluated.

6.2 Methods for Deriving Affordance Estimates

6.2.1 Probability Distribution Estimates of Affordances (PDA)

PDA take counterevidence against actions into account. Thus, their calculation is based on all known stability outcomes summarized in Table 6. There are two obvious advantages this method has when compared to point estimates. First, consider the following case: All grasp types consistently lead to the loss of an object. Only one grasp yields unstable grasp configurations most of the time. If we just considered each action's chances to result in a stable configuration, all of the grasps would be deemed equally bad. On the other hand, if the estimate was also based on the probabilities of an unstable or failed grasp, we could see that this last grasp is relatively better than the others. A second complication arises in the case of particularly unevenly spread effect CPDs for a given object. Consider the case where all actions lead to a stable grasp configuration with a probability of 50%. The other 50% are allocated to the event of losing the object for all other grasps save one. This last grasp has 50% chance to result in an unstable configuration. In case we would not consider the two additional outcomes of unstable and failed grasp configurations, it would be impossible to score this last grasp any better than the others.

To avoid these undesirable outcomes, the CPD over stability outcomes, conditioned both on action and object, is transformed into a CPD over actions, conditioned only on a given object, as shown in Equation 4.

$$P(a|o) \leftarrow \frac{\sum_{s \in S} w_s \cdot P(s|o,a)}{\sum_{a \in A} \sum_{s \in S} w_s \cdot P(s|o,a)} \tag{4}$$

where S and A are the sets of possible stability and action outcomes respectively. w_s refers to weights assigned to each stability outcome. These weights serve the purpose of increasing the likelihood of undertaking an action on an object if the observed configuration was stable and decreasing the likelihood for unstable or failing configurations. All results in this thesis were obtained with the following weights:

1. $P(S = \mathit{stable}|a, o)$ weighted with $+3$, desired outcome
2. $P(s = \mathit{unstable}|a, o)$ weighted with -1, undesired, but still not that bad
3. $P(s = \mathit{lost}|a, o)$ weighted with -4, worst case

For a given object, we calculate the weighted sum of predicted stabilities for a given action. In order to ensure a valid probability distribution with $\sum_{a \in A} P(a|o) = 1$ for a given object o, we need to keep track of the cumulative mass assigned to all actions.

This is why we need the denominator to normalize each score. An algorithm for computing the probability distributions is outlined in Listing **??** on page **??**.

Note that the proposed procedure amounts to treating $P(s|a, o)$ during the transformation phase as a pure metric, rather than as a probability distribution. It is only ensuring towards the end that the resulting scores can be considered a probability distribution.

One significant drawback of this method consists of the confidence we can have in the predicted probability distribution. This concernment is twofold: first, how powerful were the probability estimates underlying the algorithm? In fact, this is a problem that needs to be addressed one step earlier, namely in testing and evaluating the inference capabilities of a given network. The other concern is more subtle: what if the probabilities of all or of a subset of actions are about equal? Roughly said, this could express either the fact that these grasps are "equally bad" or "equally good" with respect to a given object. With the given procedure we cannot yet determine which of these cases holds. Thus, there is a need for a measure yielding information about the qualitative features of the underlying distributions. This can be implemented in various ways: one possible and computationally cheap procedure is to calculate the means and standard deviations of all possible stability outcomes (stable, unstable, lost) for either all or an interesting subset of the actions. When we then compare these measures to each other for each of the stability outcomes, it should be clear that the highest mean indicates which outcome dominates the obtained CPD.

Lastly it should be noted that an implementation of Equation 4, although conceptually easy, is not computationally cheap. This is mainly due to the need to infer the distribution over effects for all possible actions. Exact inferesnce using JTA, which is used throughout this work, becomes infeasible quickly. Thus, it is likely to become necessary to use approximate rather than exact inference algorithms when transferring this procedure to larger problem domains.

6.2.2 Methods for Point Estimates of Affordances (PEA)

There are situations in which we are not in need of a completely specified probability distribution over actions given an object, but are more interested in a quick estimate of what action might be the best to apply. In such cases it would not be economical to first compute the probability distribution $P(A|o)$ as outlined in the previous section and only then determine the action maximizing it. Thus, this section deals with point estimates of affordances (PEA), which simply means that the maximal exploitation of available evidence is traded off for computational speed. This quick estimate can be calculated using formula 5:

$$a_{aff} = arg\underset{a}{max}\ P(a|O = o, S = \text{``}stable'') \tag{5}$$

This formula amounts to pretend having already observed both the object as well as the action outcome. The action outcome was observed to be stable. Then, the action with the highest probability of explaining these two events is sought.

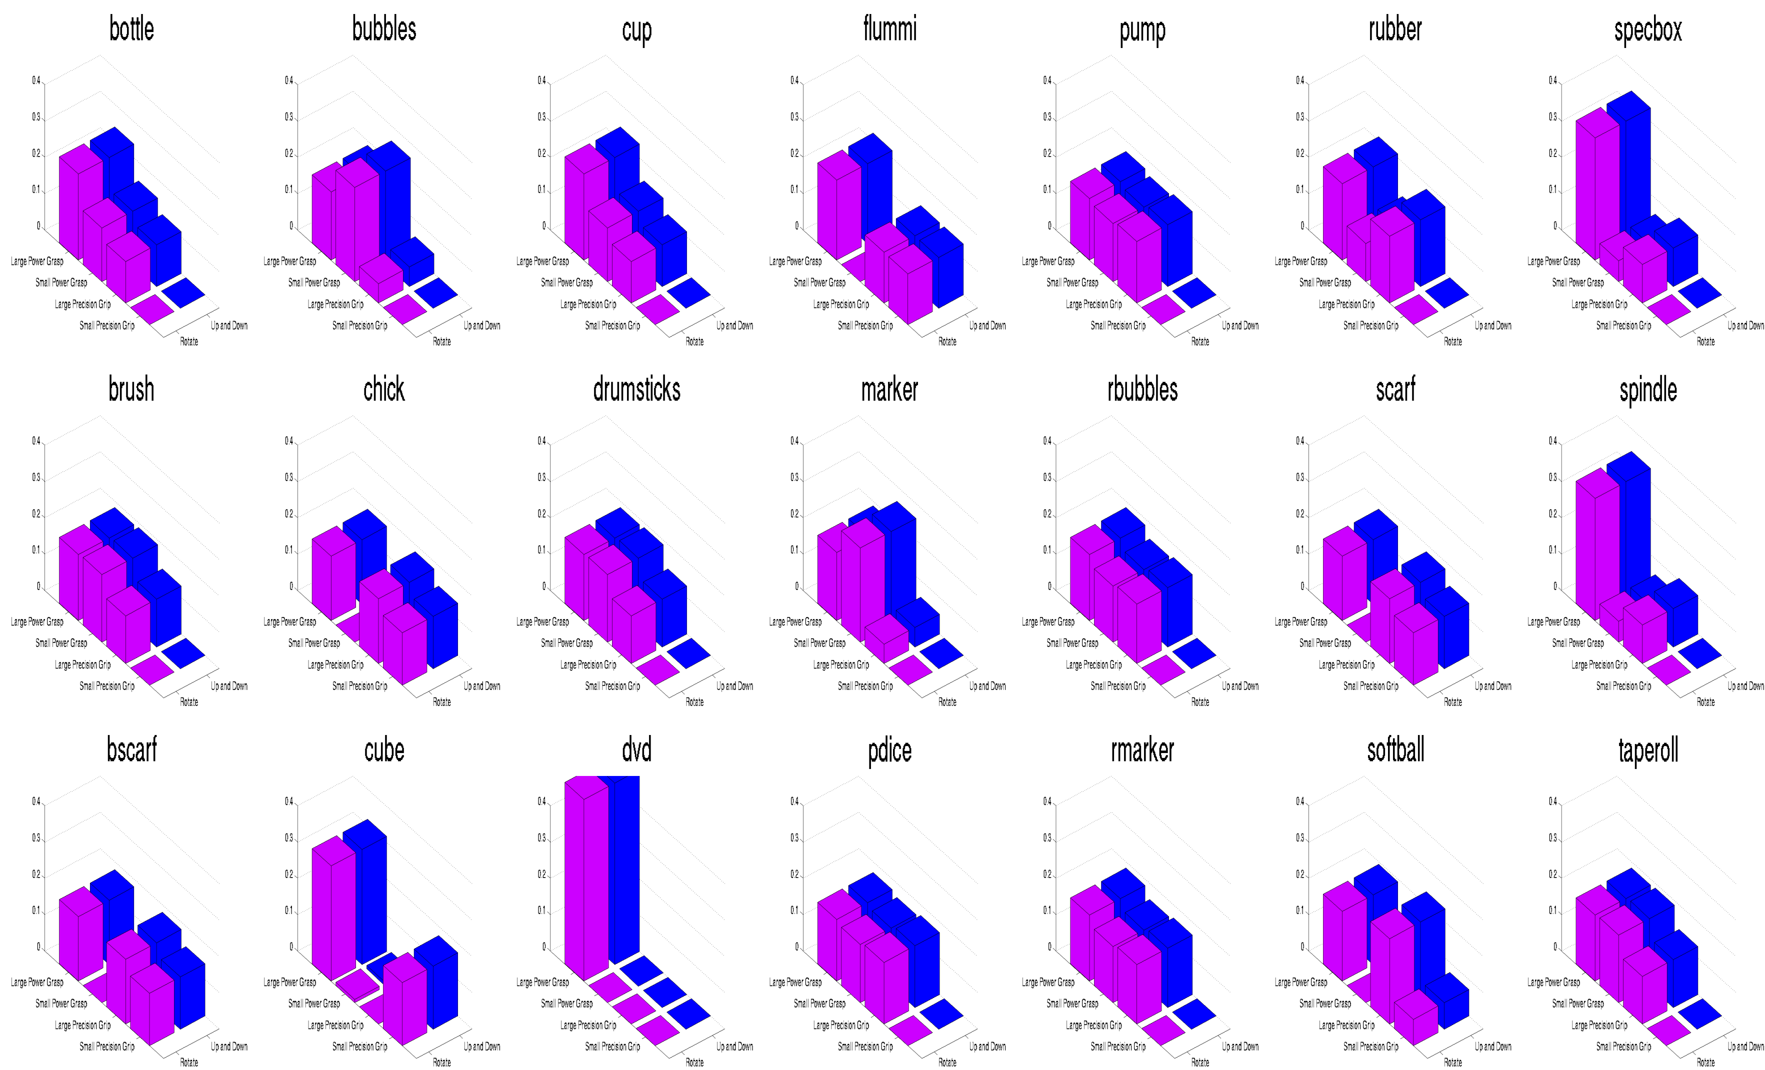

Figure 17: Distribution $P(A|o)$ as predicted by DP network

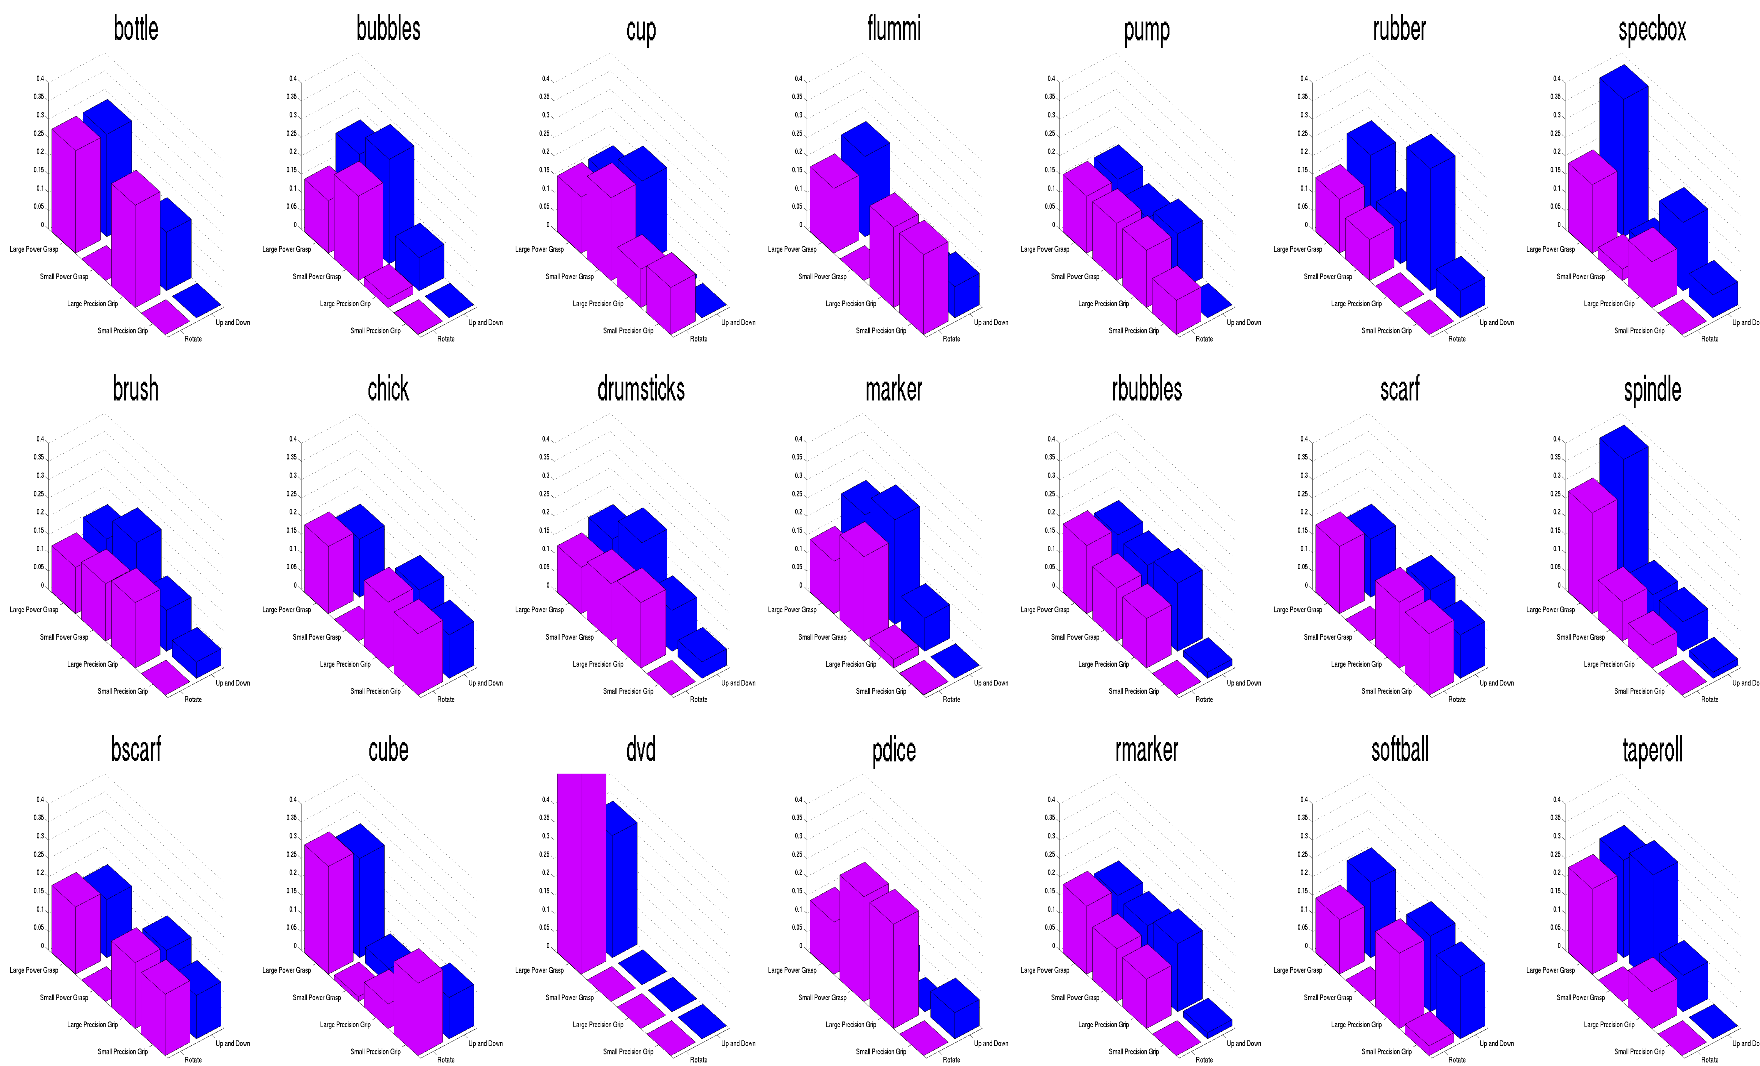

Figure 18: Distribution $P(A|o)$ as predicted by EN network

6.3 Results

6.3.1 Results obtained with PDE

Figure 17 on page 35 shows the affordance distributions over actions as computed by the DP network trained on the whole data set, Figure 18 on page 36 shows the corresponding distribution for the EN network. A distinguishing characteristic is that the affordance PDFs computed by the DP BN disregards probe as an important contribution to grasp stability and thus the the violet (rotate probing) and blue (up and down probing) columns do not differ from each other. In contrast, the EN network takes the probing action into account and thus the violet and blue columns differ.

Once more, it is difficult to establish an absolute benchmark for the computed distributions. Though, it seems apt to consider the actions recommended for the specifically tailored reference objects listed in Table 2 on page 16. For the affordance distributions of the DP network shown in Figure 17 on page 35, the respective maxima for the reference objects marker, softball and spindle match the tailored grasp: small power grasp (associated probability for this grasp type with either probing movement is ≈ 0.26), large precision grip (≈ 0.24), and large power grasp (≈ 0.34). For the last reference object, the cube, we observe that the affordance predicted by the network is a large power grasp (≈ 0.32), which is significantly better than the small precision grip (≈ 0.17) the object was intended for.

The results obtained by the EN network (Figure 18 on page 36) are similar. For fairness of comparison, the probing movement is disregarded for the EN BN, i.e. only the grasp type and the respectively higher probability of the probing movements is reported. For the objects marker (≈ 0.29) and spindle (≈ 0.38) the intended grasp types (small power and large power grasp respectively) were predicted. As was the case for the DP network, the recommended grasp for the cube was a large power grasp (≈ 0.29). Finally, the EN BN predicts a very slight advantage of a large power grasp (≈ 0.2072) over a large precision grasp (≈ 0.2055) for the object softball.

Lastly, let us illustrate the utility of additionally computing the means and STDs of stability outcomes over all actions: consider the affordance distributions obtained by the DP network for the two objects pump and brush (see Figure 17 on page 35). Both measures were only calculated for the three grasps yielding similar probabilities (large precision, large and small power grasp). Table 7 summarizes the results and also shows that the most common outcome for these three actions for the brush was a failure to grasp (≈ 0.476), whereas the pump was usually grasped successfully (≈ 0.840). Furthermore, there was more variation observed across different trials for the brush than for the pump.

	brush		pump	
Outcome	Mean	STD	Mean	STD
stable	0.235	0.077	0.840	0.000
unstable	0.289	0.172	0.120	0.030
lost	0.476	0.119	0.041	0.030

Table 7: Means and Standard Deviations of stability outcomes for the objects brush and pump obtained with the DP network

6.3.2 Results obtained with PEA

The results obtained with the point estimates method for the reference objects are shown below in Tables 8 and 9 for the DP BN and EN BN respectively. The third column gives the probability of these actions given that the outcome is set to be stable. The fourth column indicates whether the recommended grasp matches the grasp the object was intended for.

object	recommended grasp	$P(a\|e, o)$	matches?
marker	small power grasp	0.60	yes
softball	large precision grasp	0.48	yes
spindle	large power grasp	0.79	yes
cube	large power grasp	0.63	no

Table 8: Grasps recommended by PEA for DP network

The corresponding results for the EN network are shown in Table 9.

object	recommended grasp	$P(a\|e, o)$	matches?
marker	small power grasp	0.35	yes
softball	large power grasp	0.22	yes
spindle	large power grasp	0.59	no
cube	larger power grasp	0.36	no

Table 9: Grasps recommended by PEA for EN network

	DP				EN			
object	PEA		PDA		PEA		PDA	
bottle	1	0.3840	1	0.4745	1,2	0.2808	1,1; 1,2; 3,1	0.2800
brush	2	0.4523	2	0.3722	2,2	0.3161	2,2	0.2233
bscarf	3	0.2922	3	0.3577	1,1	0.1497	1,1	0.1840
bubbles	2	0.6034	2	0.5174	2,2	0.3486	2,2	0.2856
chick	3	0.2922	3	0.3577	1,1	0.1497	1,1	0.1840
cube	1	0.6312	1	0.6342	1,1	0.3628	1,1	0.2949
cup	1	0.3840	1	0.4745	2,2	0.2048	2,1;2,2	0.2262
drumsticks	2	0.4523	2	0.3722	2,2	0.3161	2,2	0.2233
dvd	1	0.2567	1	1.0000	1,2	0.1288	1,1	0.6667
flummi	1	0.4297	1	0.4402	1,2	0.2310	1,2	0.2196
marker	2	0.6034	2	0.5174	2,2	0.3486	2,2	0.2856
pdice	1	0.2917	1	0.3400	3,1	0.1677	2,1; 3,1	0.2857
pump	1	0.2917	1	0.3400	1,2	0.1637	1,1; 1,2; 3,1; 3,2	0.1563
rbubbles	1	0.3694	1	0.3636	1,1	0.1978	1,1	0.1864
rmarker	1	0.3694	1	0.3636	1,1	0.1978	1,1; 3,2	0.1864
rubber	3	0.4011	1	0.4211	3,2	0.4174	3,2	0.3333
scarf	3	0.2922	3	0.3577	1,1	0.1497	1,1	0.1840
softball	3	0.4789	3	0.4711	1,2	0.2260	3,2	0.2073
specbox	1	0.7879	1	0.6735	1,2	0.4355	1,2	0.3750
spindle	1	0.7879	1	0.6735	1,2	0.5973	1,2	0.3758
taperoll	2	0.4523	2	0.3722	2,2	0.3540	2,2	0.3000

Table 10: Affordances based on all available data

	DP				EN			
object	PEA		PDA		PEA		PDA	
bottle	2	0.3940	2	0.5979	1,2	0.1290	ALL	0.1250
brush	1	0.4669	1	0.5000	3,2	0.1292	ALL	0.1250
bscarf	3	0.2918	3	0.3428	3,1	0.1542	3,1	0.1869
bubbles	1	0.2565	ALL	0.2500	1,1	0.1286	ALL	0.1250
chick	3	0.2962	3	0.3757	3,2	0.1582	3,2	0.1826
cube	1	0.2571	ALL	0.2500	1,1	0.1290	ALL	0.1250
cup	1	0.5231	1	0.5600	1,2	0.1290	ALL	0.1250
drumsticks	1	0.4669	1	0.5000	3,2	0.1292	ALL	0.1250
dvd	1	0.2570	ALL	0.2500	1,2	0.1290	ALL	0.1250
flummi	1	0.2590	ALL	0.2500	1,1	0.1303	ALL	0.1250
marker	1	0.2565	ALL	0.2500	1,1	0.1286	ALL	0.1250
pdice	1	0.3237	1	0.3500	1,2	0.1290	ALL	0.1250
pump	2	0.2821	2	0.4000	1,2	0.1290	ALL	0.1250
rbubbles	2	0.3622	2	0.3684	1,1	0.1933	1,1; 2,1; 2,2; 3,2	0.1803
rmarker	1	0.3976	1	0.3962	1,2	0.2030	1,1; 1,2; 3,2	0.1930
rubber	1	0.2570	ALL	0.2500	1,2	0.1290	ALL	0.1250
scarf	1	0.3106	1	0.4286	1,1	0.1658	1,1	0.1970
softball	1	0.2589	ALL	0.2500	1,2	0.1311	ALL	0.1250
specbox	1	0.8984	1	0.6766	1,2	0.1290	ALL	0.1250
spindle	1	0.6141	1	0.6667	1,2	0.1287	ALL	0.1250
taperoll	2	0.4891	2	0.4126	1,2	0.1290	ALL	0.1250

Table 11: Affordances obtained with cross validation

6.3.3 Comparison of Results w.r.t. Method and Network Topology

Table 10 summarizes the actions recommended by PEA and PDA for all objects for both DP and EN topologies based on all available data. Table 11 shows the corresponding actions obtained with leave one out cross validation. The left halves list the results obtained with the DP network whereas the right halves address the results obtained with the EN network. Within each of these compartments, data is arranged as follows: the first column shows the action predicted by the point estimate (PEA) procedure along with the probability this action was attributed to in the second column. The third and fourth column hold the respective actions and probabilities for the probability distribution estimates (PDA). Note that the EN network, as seen in the right halves of the tables, can predict the grasp type as well as the probe, thus the action is an ordered pair. Since probe does not have an influence on the DP network, the action in the left halves of the tables consist only of a single integer. Please consult Table 14 on page 51 for how actions were encoded.

A quantitative dimension of the comparison between PDA and PEA is summarized in Table 12. Table 12 shows the mean and STDs of the differences between the probabilities associated with the predicted grasps by PEA and PDA. Both measures were computed using only the probabilities of those objects for which the predicted grasp was either identical or the PEA was contained in the PDA. Regardless of the network topology and what data was available, the mean of the differences approaches zero. This indicates that there are no systematically larger probabilities entailed by either PEA or PDA. On the other hand, the standard deviations in cases where all data was used are larger than those obtained with cross validation. This difference, however, can be solely attributed to the object dvd, which yields particularly different probabilities whether PEA (≈ 0.2567) or PDA (1.000) was used. When disregarding the dvd in the calculations, STDs approach 0.07 for both networks.

Qualitatively, consider the type of grasps recommended by PDA and PEA when all data was used, as shown in Table 10. The DP network yields the same predic-

	Mean	Std
DP all data	-0.028	0.178
EN all data	0.004	0.141
DP cross validation	-0.019	0.081
EN cross validation	0.000	0.013

Table 12: Mean and standard deviations of differences between predictions by PEA and PDA

tions for grasp type whether PDA or PEA is used for 20 out of 21 objects. The EN network performs a little worse with identical predictions for 15 objects. For another five objects, the PEA is contained in the set of the most likely actions as predicted by PDA. These preliminary observations contrast sharply with the affordance estimates obtained through leave one out cross validation: Table 11 illustrates that the PDA method produces more uniform predictions than the PEA counterpart for both network types. In numbers, eight uniform predictions for the DP network and 16 for the RN network are observed. To a certain extent, this was to be expected based on the inference results of the respective networks as discussed in chapter 5.

Furthermore, in cases where PDA predicted non-uniform distributions, the corresponding PEA did not resort as strongly to random baseline predictions for the DP network as they did for the EN network. To illustrate this, the means of the probabilities of the grasps obtained through PEA[17] were calculated for both topologies. For the EN network, the mean equals 0.1761, which amounts to an increase in certainty of 0.4088% over the random baseline of 0.125. The mean for the DP network is 0.4369, which amounts to an increase in certainty of 0.7476% over the baseline of 0.25[18]. Correspondingly, when all data was used, the amount of increase in certainty was 0.76% for the DP network and 121.88% for the EN network.

6.4 Discussion of Affordance Modeling Efforts

Considering the results for computing affordances using either point estimates (PEA) or probability distribution estimates (PDA) separately, little else can be said than that both methods did relatively well at predicting the most stable grasps for the reference objects. On a side note, the fact that both networks predicted a large power, rather than a small power grasp for the cube, cannot be considered counterevidence for the models' affordance prediction capabilities. Rather, it was actually the case [19] that a large power grasp resulted in a relatively more stable grasp.

Overall, it seems that network topology was a more influential factor than the method chosen for computing affordances. One result supporting this is that there is no systematic trend that for a given network either the PDA or PEA produces higher estimates than the other method (see Table 12). On the other hand, the different generalization capabilities of the DP and EN topologies are subtle: when all data is used, the performance of either network is fine. However, it seems that the engineered

[17]Only PEA recommendations congruent with the PDA results were used.

[18]To avoid confusion, the respective baselines differ from each other since the RV action can take on eight values (two probes and four grasps) for the EN network and only four for the DP network.

[19]As the reader may verify from the visualization of the MLE distribution for the cube shown in Figure 14, page 26

network is almost doing too good in the sense that the probabilities it associates with it's predictions are astonishing 120 percent above the network's random baseline. This, paired with the fact that the same number drops to 40 percent when cross validation is used, strongly suggests that the engineered network is overfitting the available data. This also explains why the PDA for the EN network is biased to produce uniform distributions. The DP network improves in both cases, that is when all data is used or cross validation is applied, by about 75 percent with respect to its baseline. This constant rate of improvement suggests that the available data is well fitted by the DP network. However, it is necessary to point out that the relatively better inference performance of the DP network when compared to the EN network is a trade-off for the less specific recommendations of actions the DP network yields. Specifically, the DP cannot distinguish between different probing movements and their influence on the observed grasp stability.

We have seen that both estimates for the most likely action, whether based on computing a complete probability distribution over actions first or simply approximating it by taking a point estimate, are suitable measurements for affordances. Indeed, for our specific data set we observed that the predictions for the most likely action are the same for most objects using the DP network. In the particular case of this work, the good performance of PEA is not too surprising, since it amounts to being able to guess the probability distribution of a RV that can take on only three values given the probability of one of them. However, it is not guaranteed that the heuristic of guessing these two remaining values works well. The only object in our data set which clearly exemplifies this is the dvd. In the case of the DP network trained with all data, we see that both PEA and PDA yield the identical estimates that a large precision grasp should be applied to the dvd. However, the certainty with which PEA recommends this action (0.2567) is much below that of PDA (1.000). These values can be interpreted as PEA recommending further exploration to determine the best action, whereas the PD estimate comforts us with having found a good action with a probability of 1. The basis of this difference can be seen in the MLE estimation of outcomes for the combination of each grasp with the dvd, as has been illustrated in 14 and 16 on pages 26 and 28. The MLE estimates show clearly that all grasps lead to a loss of the dvd, except the large power grasp which yields unstable results most of the time. Thus, since the large power grasp is still the best among many bad choices, there is evidence it should be the preferred grasp type. This is the result returned by PDA. Another way of reasoning is that we actually do not want to take any of these grasps, given that we cannot reach the goal of a stable grasp in any case. This is what is computed by the PEA estimate. Therefore, the two measures can be seen as complementing each other, with the PDA measure providing a ground on which we can choose sensible actions although our action set is limited, and PEA pointing out that indeed none of the grasps seems perfect.

To conclude we can say that if we are willing to make the assumption that at least one grasp primitive is more or less suitable for any object encountered, PEA seems to be a suitable method. This holds especially when the goal is to quickly compute an applicable result. In general, however, PDA can be expected to yield more reliable results. It is also needed where a valid probability distribution over actions is required. This is the case, for example, when affordances are used as input for other computations, such as in the mirror neuron model by Metta et al.

7 Conclusions and Future Directions

7.1 Summary of Results

In the preceding sections, an effort was made to attain the goals set forth in section 1.3. First, it was aimed at replicating the results of Montesano et al. in the domain of grasping and for a set of more variable objects than the ones used in the original work. Taking a reductionist approach, this replication task decomposes to learning the structure of a Bayesian Network and to demonstrating that this network has sensible inference capabilities.

Overall, the task of structure learning was completed successfully: all of the structure learning algorithms applied, be it DP, DPMCMC or K2, managed to find dependencies between some set of object and action features towards the effect of stability. As exemplified in Section 4.3.4, these results can only be guaranteed if interventional data and, optionally, temporal layering information is available to the structure learning algorithm. Thus, providing this information is essential since without a sensible network topology it cannot be expected that the prediction of future events will function correctly. But still if providing interventional data to the algorithms, we found differences in the inference performance of different topologies. These can be summarized as follows: DP and DPMCMC yielded satisfactory results for both the inference process and the computation of affordances. The greedy K2 algorithm yielded poor results throughout. The EN network ranked second in performance after the DP network. There were indications that the sparse connectivity of the K2 BN led to underfitting of the data while the highly connected EN topology entailed overfitting of the data.

The poor performance of the K2 algorithm in particular was unexpected since this algorithm provided good results for Montesano et al. This could be attributed to the higher number of random variables assigned to effects in their experimental layouts. In summary, the results by Montesano et al. could be replicated. The most distinguishing difference between the work of Montesano et al. and this thesis is that all information was acquired in an unsupervised manner in the former. In order to do the same in the domain of grasping, at least a satisfying solution to the problem of unsupervised acquisition of stability would be necessary.

The second goal of this thesis was to demonstrate a technique for arriving at the affordance representation introduced by Metta et al. based on the acquired networks. The essential characteristics of these probability distribution estimates of affordances were to consider stability as the implicit goal of a grasping action as well as treating the inferred effect probabilities as metrics during the transformation process. This goal was attained and the results are extensively discussed in section 6.4. Although the proposed technique can be generalized to multiple variables, it is unlikely that there will be a practical application of the attained distributions unless its computational costs are reduced. However, the use of point estimates, requiring only a single inference operation might be a feasible and more practical alternative in some contexts.

7.2 Discussion

I would like to remark that the framework of Bayesian Networks proved itself very useful for affordance modeling. This was partly due to the fact that they are transparent models, i.e. the internal probability distributions of nodes can be inspected and compared against the gathered evidence at any point in time. Furthermore, the fact that Bayesian Networks are well studied graphical models and numerous algorithms along with powerful implementations exist, makes them suitable for more exhaustive experiments. A significant drawback, however, is the static nature of Bayesian Networks which makes it difficult to add or prune edges during learning. To accommodate this attribute, it seems advisable to make use of ensemble techniques to keep more than one network at a time. Then, it is likely to become necessary to reevaluate and partially replace this set on a rolling basis.

An interesting aspect of modeling affordances using Bayesian Networks is that it is possible to seamlessly integrate a top down motivational influence on grasp choice. Obviously, not only during the learning phase the motivation of the action is a crucial factor influencing what action is executed: whether grasping a glass in order to smash it onto a wall or in order to drink out of it is a major difference. To mimic this on a lower level, two probing movements exerting different forces on the held object were designed. The rationale behind this was, that certain objects would withstand the vertical forces better than rotational forces and vice versa. For example, a long, elongated object such as the drumsticks, would be expected to be less stable when rotated rather than lifted up. The Bayesian Network performing best throughout the experiments, however, assumed there to be no influence of the probing action on the observed stability. This might be due to the fact that there were only two components constituting an action and the probing actions were too similar. Therefore, we could not directly demonstrate the flexibility of the chosen formalization approach with respect to motivational influences. As a foretaste of how the different components of an action can influence object stability could look like, we can examine the affordance distribution $P(a|o)$ as computed for the EN network and shown in 18 on page 36.

There we can see that some objects, e.g. the specbox, resisted much better to the up and down movement rather than to the rotational movement.

A major limitation of the demonstrated technique to compute affordances is that it made use of human knowledge in that it was asserted that stability is the goal of a grasping action. Only through this assumption it was possible to approximate the "usual" action for an object. This does in no way solve the general problem of how an autonomous agent should know or determine what the desired effects for different objects are: as a simple example, consider the water pump illustrated in Section 1. For this object, it seems intuitive to assume water flow as the goal with respect to which actions should be evaluated. This problem of associating objects with desired effects is likely to require a completely separate mechanism.

7.3 Future Directions

Owing to the exploratory nature of the work presented, there are many possibilities for future research. I would like to go into more detail for three of these possibilities, which I consider to be at one hand viable and at the other hand interesting to explore.

First, I think that representing actions themselves as Bayesian Networks is a possibility that should be investigated. This topic was already tentatively dealt with by Montesano et al. For example, they demonstrated that treating the height of the wrist at the moment of grasping itself as a random variable, BALTAZAR could adapt this free parameter so to make this grasp a successful one. Since structure learning as well as inference in Bayesian Networks quickly becomes complex as the number of variables increases, it is certainly not advisable to assign a random variable to each degree of freedom of the given robotic platform. Rather, an intermediate level of description is needed. Ideally, this level would decompose actions into meaningful subunits, such as the height of the grasp, the proximal closure angle of the fingers or the orientation of the wrist. Whether it would be more practical to consider the BN finally constituting an action as a closed subsystem or to integrate it seamlessly into a BN also operating on the level of object features and effects, needs to be researched empirically. Another side note on this is that at first sight, it seems that the idea of composing actions from specific motor features parallels the use of motor features in the model of the mirror neuron system by Metta et al.

The second direction, which I consider to be of fundamental importance, is to increase the agent's autonomy based on better availability and better exploitation of sensory information. Specifically, the agent should be enabled to make better use of visual as well as tactile or proprioceptive stimulation. In the domain of vision, object features could be observed directly and would not have to be provided from the outside. For example, Chinellato and colleagues presented an interesting technique based on object exploration which is suitable for object shape detection [37]. A better exploitation of visual information is also important to equip the robot with the capability of performing visually guided or visually elicited reaching, that is, being able to locate an object and its orientation in space[20]. Furthermore, it might also be possible to integrate the capabilities for biological motion detection, as they were e.g. described in [42], to visually aid the process of action recognition, instead of merely relying on the observed object and effects. Also, the gain of the obvious factor of tactile and proprioceptive data in reaching is worth exploring, which might enable the agent of acquiring a more direct quantity for the stability measure. Also, the tactile information could aid in specifying object features which are not easily extracted visually, e.g. rigidity or texture [36]. Finally, tactile feedback information could be used to enable the robot to adapt grasp parameters online.

Third, an affordance learning system such as the one outlined here could profit considerably from motivational modulation and "curious" behavior. To see this, consider that the implicit goal of stability can be interpreted as an equivalent to the "joy of graspring" observed in children [43], thus treating stability as a motivational factor. Children, unlike robots, are innately eager to learn about their surroundings and show contentment when they manage to manipulate it. Of course, this curiosity develops dynamically, and once the child realizes how it can reliably reproduce a certain manipulation, this manipulation will tend to become boring. Paralleling these observations, it seems indeed that in order to put the approach to affordance learning put forward in this work to practical usage, it needs to be complemented with a motivational system aiding the robot in seeking out new stimulation, be it visual or ultimately tactile or proprioceptive. Such motivational or curiosity modules recently gained more interest in the robotics community, so to combine these two methods would constitute an interesting direction to explore.

7.4 Conclusion

This thesis introduced an affordance based approach to enable robots to grasp objects in their surroundings successfully. It stepped into the gap between the methodology demonstrated by Montesano et al. and the theoretical model of the mirror neuron system introduced by Metta et al., specifically filling in the affordance prior. At a certain level of abstraction, where a lot of problems arising in the domain of early

[20]For telling experiments on the importance of this visual "preexploration", see [40] and [41]

sensory processing are already solved satisfactorily, it seems that the probabilistic approach taken in this thesis is promising both in that it can adequately address the problem of predicting effects from everyday experience and that the nature of its representation makes it to use this knowledge in actual robot control. In future research, I see the place of Bayesian Networks at an intermediate complexity level, fed from lower levels with sensory information and the capabilities for motor action planning and modulated through attentional processes. Indeed, the possibility to integrate such top down processes seamlessly is an exciting possibility.

Overall, this thesis can be considered another building block in the effort to create more autonomous robots, especially facilitating interaction with their surroundings. With interaction I have in mind a threefold implication: first, the interaction with people, second, the interaction with objects and surfaces, and lastly, the important capability of predicting the future outcome of events, which already seems to come so easily to infants.

8 Appendix

8.1 Abbreviations

BN	Bayesian Network
BDeu	Bayesian Dirichlet uniform equivalent
CPD	Conditional Probability Distribution
D_{KL}	KL Divergence
DAG	Directed Acyclic Graph
DP	Dynamic Programming
DPMCMC	Dynamic Programming combined with Markov Chain Monte Carlo
EN	Engineered network
JTA	Junction Tree Algorithm
MAP	Maximum A Posteriori Estimate
MCMC	Markov Chain Monte Carlo
MLE	Maximum Likelihood Estimation
PEA	Point Estimate of Affordances
PDA	Probability Distribution Estimate of Affordances
PDF	Probability Density Function
RV	Random Variable
STD	Standard Deviation

8.2 Software

All analysis in this thesis was performed with software based on the Bayesian Network Toolbox and the Bayesian DAG Learning Toolbox for MATLAB, both of which were developed by Kevin Murphy[44]. The software is available at http://www.cs.ubc.ca/ murphyk/Software/.

8.3 Object Feature Encoding

Feature Name	Shape	Height	Rigidity	Texture	Color
Arity	4	2	3	3	3
Arity Labels	1: long, narrow	1: short	1: hard	1: smooth	1: bright
	2: long, wide	2: tall	2: normal	2: normal	2: other
	3: short, narrow		3: soft	3: textured	3: dark
	4: short, wide				

Table 13: Object Feature Encoding

8.4 Grasp and Probe Encoding

Grasp 1	Large Power Grasp
Grasp 2	Small Power Grasp
Grasp 3	Large Precision Grasp
Grasp 4	Small Precision Grasp
Probe 1	Rotation of the Forearm
Probe 2	Up and Down Movement of whole Arm

Table 14: Grasp and Probe Encoding

8.5 Joint Parameters of Grasp Types

Joint	LPow	SPow	LPrec	SPrec
Thumb opposition	17	12	88	68
Thumb proximal flexion/extension	54	9	34	32
Thumb distal flexion	44	9	35	22
Index proximal flexion/extension	58	60	35	54
Index distal flexion	51	81	50	10
Middle proximal flexion/extension	59	89	30	72
Middle distal flexion	50	78	50	10
Ring & little finger flexion	109	120	107	70

Table 15: Joint Parameters for iCub's Hand for all primitive grasping actions. All values are given in Degrees.

8.6 Overview of used Objects

Object Name	Shape	Height	Rigidity	Texture	Color
bottle	2	2	2	1	2
brush	1	1	1	2	1
bscarf	4	2	3	3	2
bubbles	3	1	1	1	3
chick	4	2	3	3	1
cube	3	1	2	1	1
cup	4	2	2	1	2
drumsticks	1	1	1	2	1
dvd	2	1	2	1	3
flummi	3	2	1	3	2
marker	3	1	1	1	3
pdice	3	2	2	3	2
pump	2	2	2	2	2
rbubbles	3	1	1	3	3
rmarker	3	1	1	3	3
rubber	3	1	3	3	2
scarf	4	2	3	3	3
softball	4	2	3	2	1
specbox	4	2	1	2	3
spindle	2	2	1	2	1
taperoll	3	1	1	2	1

Table 16: Summary of objects used and their respective feature values.

Figure 19: Pictures of Objects used. The ordering of the objects corresponds to the alphabetical ordering of their names, therefore these pictures correspond to the listing on the previous page if read line by line.

References

[1] L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor, "Learning object affordances: From sensory-motor coordination to imitation," *IEEE Transactions on Robotics*, vol. 24, no. 1, pp. 15–26, 2008.

[2] G. Metta, G. Sandini, L. Natale, L. Craighero, and L. Fadiga, "Understanding mirror neurons: a bio-robotic approach," *Interaction studies*, vol. 7, no. 2, pp. 197–232, 2006.

[3] J. Gibson, *The Ecological Approach to Visual Perception.* Houghton Mifflin, 1979.

[4] J. J. Gibson, "The concept of affordances," *Perceiving, acting, and knowing*, p. 6782, 1977.

[5] E. J. Gibson and A. D. Pick, *An ecological approach to perceptual learning and development.* Oxford University Press, USA, 2000.

[6] W. H. Warren, "Perceiving affordances: Visual guidance of stair climbing," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 10, no. 5, pp. 683–703, 1984.

[7] S. Ishak, K. E. Adolph, and G. C. Lin, "Perceiving affordances for fitting through apertures," *Journal of experimental psychology. Human perception and performance*, vol. 34, no. 6, p. 1501, 2008.

[8] J. R. Flanagan and R. S. Johansson, "Action plans used in action observation," *Nature*, vol. 424, no. 6950, pp. 769–771, 2003.

[9] A. L. van der Meer, "Keeping the arm in the limelight: advanced visual control of arm movements in neonates," *European Journal of Paediatric Neurology*, vol. 1, no. 4, pp. 103–108, 1997.

[10] C. von Hofsten, "An action perspective on motor development," *Trends in cognitive sciences*, vol. 8, no. 6, pp. 266–272, 2004.

[11] G. Rizzolatti and G. Luppino, "The cortical motor system," *Neuron*, vol. 31, no. 6, pp. 889–901, 2001.

[12] L. Fadiga, L. Fogassi, V. Gallese, and G. Rizzolatti, "Visuomotor neurons: Ambiguity of the discharge or motorperception?," *International Journal of Psychophysiology*, vol. 35, no. 2-3, pp. 165–177, 2000.

[13] G. Pellegrino, L. Fadiga, L. Fogassi, V. Gallese, and G. Rizzolatti, "Understanding motor events: a neurophysiological study," *Experimental brain research*, vol. 91, no. 1, pp. 176–180, 1992.

[14] M. Jeannerod, M. A. Arbib, G. Rizzolatti, and H. Sakata, "Grasping objects: the cortical mechanisms of visuomotor transformation," *Trends in Neurosciences*, vol. 18, no. 7, pp. 314–320, 1995.

[15] A. H. Fagg and M. A. Arbib, "Modeling parietal-premotor interactions in primate control of grasping," *Neural Networks*, vol. 11, no. 7-8, pp. 1277–1303, 1998.

[16] E. Charniak, "Bayesian networks without tears," *AI magazine*, vol. 12, no. 4, pp. 50–63, 1991.

[17] C. M. Bishop, *Pattern recognition and machine learning.* Springer New York., 2006.

[18] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, 1988.

[19] D. Heckerman, "A tutorial on learning with bayesian networks," *Learning in graphical models*, pp. 301–354, 1998.

[20] R. Robinson, *Counting labeled acyclic digraphs.* Academic Press, 1973.

[21] D. M. Chickering, "Learning bayesian networks is NP-complete," *Learning from data: Artificial intelligence and statistics v*, p. 121130, 1996.

[22] G. F. Cooper and E. Herskovits, "A bayesian method for the induction of probabilistic networks from data," *Machine learning*, vol. 9, no. 4, pp. 309–347, 1992.

[23] M. Koivisto and K. Sood, "Exact bayesian structure discovery in bayesian networks," *The Journal of Machine Learning Research*, vol. 5, pp. 549–573, 2004.

[24] M. Koivisto, "Advances in exact bayesian structure discovery in bayesian networks," in *Proc. of the 22nd Annual Conf. on Uncertainty in Artificial Intelligence*, 2006.

[25] D. Eaton and K. Murphy, "Bayesian structure learning using dynamic programming and MCMC," *Proceedings of the Twenty-Third Confererence on Uncertainty in Artificial Intelligence (UAI 2007)*, 2007.

[26] T. Silander and P. Myllymaki, "A simple approach for finding the globally optimal bayesian network structure," in *Proceedings of the Twenty Second Conference on Uncertainty in Artificial Intelligence (UAI 2006),*, 2006.

[27] D. Heckerman, D. Geiger, and D. M. Chickering, "Learning bayesian networks: The combination of knowledge and statistical data," *Machine learning*, vol. 20, no. 3, pp. 197–243, 1995.

[28] F. V. Jensen, K. G. Olesen, and S. K. Andersen, "An algebra of bayesian belief universes for knowledge-based systems," *Networks*, vol. 20, no. 5, 1990.

[29] S. L. Lauritzen and D. J. Spiegelhalter, "Local computations with probabilities on graphical structures and their application to expert systems," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 50, no. 2, pp. 157–224, 1988.

[30] H. Guo and W. Hsu, "A survey of algorithms for real-time bayesian network inference," in *Joint AAAI-02/KDD-02/UAI-02 workshop on Real-Time Decision Support and Diagnosis Systems, Edmonton, Alberta, Canada*, 2002.

[31] L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor, "Affordances, development and imitation," *IEEE - International Conference on Development and Learning (ICDL 2007), London, UK*, 2007.

[32] P. Fitzpatrick, G. Metta, L. Natale, S. Rao, and G. Sandini, "Learning about objects through action - initial steps towards artificial cognition," in *IEEE International Conference on Robotics and Automation, 2003. Proceedings. ICRA 2003*, vol. 3, 2003.

[33] G. Sandini, G. Metta, and D. Vernon, "The iCub cognitive humanoid robot: An open-system research platform for enactive cognition," *Lecture Notes in Computer Science*, vol. 4850, p. 358, 2007.

[34] J. R. Napier, "The prehensile movements of the human hand," *Journal of Bone & Joint Surgery, British Volume*, vol. 38, no. 4, pp. 902–913, 1956.

[35] A. Bicchi, "On the closure properties of robotic grasping," *The International Journal of Robotics Research*, vol. 14, no. 4, p. 319, 1995.

[36] L. Jamone, G. Metta, F. Nori, and G. Sandini, "James: A humanoid robot acting over an unstructured world," in *Humanoid Robots, 2006 6th IEEE-RAS International Conference on*, pp. 143–150, 2006.

[37] E. Chinellato, G. Recatala, A. P. del Pobil, Y. Mezouar, and P. Martinet, "3D grasp synthesis based on object exploration," in *Proceedings of the IEEE international conference on robotics and biomimetics*, pp. 1065–1070, 2006.

[38] I. Ebert-Uphoff, "The linkstrength package for bnt (version 0.1),"

[39] I. Ebert-Uphoff, "Measuring connection strengths and link strengths in discrete bayesian networks," *Woodruff School of Mechanical Engineering, Atlanta, GA, Tech. Rep*, 2006.

[40] M. E. McCarty, R. K. Clifton, D. H. Ashmead, P. Lee, and N. Goubet, "How infants use vision for grasping objects," *Child development*, pp. 973–987, 2001.

[41] R. Held and J. A. Bauer, "Visually guided reaching in infant monkeys after restricted rearing," *Science*, vol. 155, no. 3763, pp. 718–720, 1967.

[42] M. A. Giese and T. Poggio, "Neural mechanisms for the recognition of biological movements," *Nature Reviews Neuroscience*, vol. 4, no. 3, pp. 179–192, 2003.

[43] E. Oztop, *Modeling the mirror: Grasp learning and action recognition.* PhD thesis, University of Southern California, 2002.

[44] K. P. Murphy, "The bayes net toolbox for matlab," *Computing science and statistics*, 2001.

Affirmation

Hereby I confirm that I wrote this thesis independently and that I have not made use of any resources or means other than those indicated.

Alice Ellmer, Los Angeles, September 2009